



HAL
open science

Experimental investigation of microstructural changes in deforming granular media using x-ray tomography

Edward Ando

► **To cite this version:**

Edward Ando. Experimental investigation of microstructural changes in deforming granular media using x-ray tomography. Mechanics [physics.med-ph]. Université de Grenoble, 2013. English. NNT : 2013GRENI097 . tel-01144326

HAL Id: tel-01144326

<https://theses.hal.science/tel-01144326>

Submitted on 21 Apr 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE

Pour obtenir le grade de

DOCTEUR DE L'UNIVERSITÉ DE GRENOBLE

Spécialité: **Matériaux, Mécanique, Génie Civil, Electrochimie**

Arrêté ministériel: 7 août 2006

Présentée par

EDWARD CARLO GIORGIO ANDÒ

Thèse dirigée par **Jacques DESRUES** et

codirigée par **Gioacchino VIGGIANI** et **Stephen A. HALL**

préparée au sein du **Laboratoire 3SR**

dans l'**École Doctorale Ingénierie - Matériaux Mécanique Environnement
Énergétique Procédés Production (I-MEP²)**

Experimental investigation of micro- structural changes in deforming granular media using x-ray tomography

Thèse soutenue publiquement le **05.02.2013**,
devant le jury composé de:

Dr. Michel BORNERT

Chercheur, École des Ponts ParisTech, Membre

Dr. Jacques DESRUES

Directeur de Recherche, CNRS UMR 5521, Directeur de Thèse

Prof. Itai EINAV

Professeur, University College London, Rapporteur

Dr. Stephen A. HALL

Associate Professor, University of Lund, Directeur de Thèse

Univ. Prof. Dipl.-Ing. Dr. Dimitrios KOLYMBAS

Professeur, Universität Innsbruck, Membre

Dr. Jean-Noël ROUX

Ing. en Chef IFSTTAR, Laboratoire Navier - Champs sur Marne, Membre

Dr. Stéphane ROUX

Directeur, de Recherche LMT-Cachan, Rapporteur

Prof. Gioacchino VIGGIANI

Professeur, Université Joseph Fourier, Directeur de Thèse



Thanks

Abstract thanks

Thanks to the strange and amazing country of France for putting me up and putting up with me.

Thanks to Laboratoire 3SR for giving me a super place to work.

Thanks to the Ministère de l'Enseignement Supérieur et de la Recherche for providing funding for my PhD.

Thanks to the powers that be for allowing me to teach during my PhD, a fascinating experience which I have enjoyed a lot.

Real people

Thanks to Carole, Sheherazade and Augustine for being the friendly face of the infamous French bureaucracy. Thanks to Rémi for help, discussions, and freedom to use computers as necessary in the lab. Not everyone is blessed with such a friendly and efficient IT guru! Thanks also to Pascal Charrier who has helped considerably with the technical side of this work.

Thanks to all my friends for existing (come on! You're really looking for your name in a PhD thesis?).

Thanks specifically to colleagues from all domains for the fascinating and sometimes overwhelming discussions we have shared.

Thanks to those with whom I have interacted in a more formal context of scientific cooperation, I have learned a lot from all the collaborations undertaken in this work, and can only hope to keep that they keep going!

Deep thanks to Vittorio, who won't be able to read this, but who is one of the principal reasons for which I have had the good fortune to come to Grenoble to work in this extraordinary place with these extraordinary people. You are one of the people who will be most missed on the day of the defence of my PhD.

Some of the most amazing people I know

I am not going to thank my parents: I owe almost everything I am to them, their total love is fundamental to me and they are profoundly involved in this work indirectly. However, their direct involvement in this work is minimal, so this is not the place to thank them.

I think the supervisors of my work – The Englishman, the Frenchman and the Italian – need to be thanked first of all, all together. Just the *nationalities* of this triumvirate correspond so perfectly with my background and with where I feel at home. It has been a continuous source of pleasure to be supervised and guided by this mad team, both at the top level as well as in the details. Your respect for my inability or slow speed in some things, and yet your continued interest and availability has made this a great experience, and I know not all PhD students are blessed with such great supervision. What continues to surprise me is the ability of each one of you – each with his experience (although you can all pretty much do everything in this domain) – to really care and understand the finer details of the work I have done.

The Englishman first, how great to have at least one mother tongue English speaker in Grenoble with whom to converse... that is until you left for Sweden! Your departure from the lab was a real pity for me (and I know I'm not alone in thinking this). Your supervision has really been very well suited to my way of working, always keeping me interested. Your work, and your way of working has also been a great source of inspiration for me – I can only hope to be able to do such interesting things in the future. The time spent together in Sweden recently was extremely fun and productive and I hope our plan to do a repeat visit will go ahead with great things to come! Thanks Steve, thanks really a lot.

Jacques : (with French punctuation) malgré le fait que tu ais été très occupé par la direction du laboratoire, tu m'as accordé ton temps précieux avec patience et intérêt pour mon travail. Les discussions sur notre sujet m'ont ravi et fasciné même si nos niveaux de compréhension peuvent, peut être, se comparer sur une échelle log! Tu m'as appris vraiment beaucoup de choses, mais je pense que la chose la plus importante est ta méthode (par excellence scientifique) de patiemment et paisiblement tout remettre en cause, pour mieux comprendre. Si je reste à Grenoble c'est aussi dans l'espoir qu'une fois fini la direction du laboratoire on pourra travailler ensemble, sur les nombreuses idées brillantes que tu me lances souvent!

It would be rude to talk in Italian, so Cino, thanks. Thanks for the time you've spent with my daily questions, thanks particularly for your patience with my hangups and stupidities. You have spent the most amount of time on this work with me, and your constant attention has shown me with how much love you approach your students. I am in awe of your way of working and your knowledge of all parts of our field. You have seen, and have principally been responsible for my gradual change from "I will leave Grenoble and go into industry if I ever get my PhD" to the current state of affairs: staying in Grenoble and working with you in this great team. If this were a novel I would probably have replaced, after deliberation, all these thanks simply by: *with Cino and for Cino.*

Note on Software

Where possible in this work Open Source software has always been preferred, both on an Operating System front, as well as for the software used for acquiring, treating and analysing data. Unfortunately in some cases, closed source operating systems and software have been used in this work, but on almost all fronts recent developments show encouraging trends in the direction of being able to replace these black boxes.

In this research work PYTHON has been used extensively, with other tools such as ImageJ (Fiji) and Gnuplot also being used widely.

Where it been considered useful, software or algorithms have been contributed back to the community.

Abstract in English

This doctoral thesis presents an experimental investigation into the mechanics of granular media. The novelty that this work brings is that the specimens of sand tested in this work are systematically and non-destructively imaged using x-ray tomography. Sample size is considerably reduced from standard (specimens measure approximately 22 mm height by 11 mm diameter), allowing entire specimens to be scanned at a sufficiently high resolution to identify all the grains (more than fifty thousand) in each specimen.

A campaign of triaxial compression tests has been run on a series of three different natural sands with different grain shapes (Hostun sand, Ottawa sand and Caicos ooids – all prepared at relatively dense initial states), and tested at 100 or 300 kPa cell pressure. In each test around 15 x-ray scans are performed. In the 3D images resulting from the reconstruction of the x-ray scans performed, grains are identified each state using a standard watershed algorithm. Starting from these discretised data, techniques are developed in order characterise grain-to-grain contacts, as well as to measure the kinematics of all the identified grains between imaged states. Grain kinematics are measured with two specifically-developed tools: “ID-Track” to track grains yielding their displacements, and a discrete image correlation technique to measure grain rotations.

Grain-scale measurements are reported in detail for one test, and are then compared to tests in different conditions, in order to highlight the micro-mechanisms responsible for the observed macroscopic behaviour. This comparison highlights some important micro-scale mechanisms such as the increasing rotational frustration of more angular grains when the sample’s deformation is concentrated in a fully developed shear band; this is used to explain to some extent the higher value of their residual stress for these materials. Signs of localised deformation are seen to occur well before the peak in many samples, and complex patterns of rotating grains (which match a local, grain-based measurement of strain) are noticed around the peak of each sample’s response.

Résumé en Français

Cette thèse présente un travail expérimental dans le domaine de la mécanique des milieux granulaires. Ce travail introduit une nouveauté fondamentale: durant la déformation d'échantillons de sable en compression triaxiale classique, leur micro-structure est enregistrée par tomographie à rayons-x en environ 15 étapes de déformation différentes.

Afin que tous les grains d'un échantillon soient individuellement visibles dans les images de tomographie, les échantillons sont considérablement réduits pour qu'ils mesurent 11 mm en diamètre et 22 mm en hauteur. Cela permet d'identifier et d'individualiser tous les plus de cinquante mille grains d'un échantillon.

Dans ce travail expérimental, une série d'essais triaxiaux a été réalisée sur trois sables naturels (le sable d'Hostun, le sable d'Ottawa et des Caicos ooids, tous à partir des états initiaux relativement denses), à deux valeurs de pression de confinement différents (100 et 300 kPa).

Dans les images 3D résultantes de la reconstruction des acquisitions tomographiques réalisées au cours de chaque essai, les grains sont identifiés dans chaque état en utilisant un algorithme de type watershed (ligne de partage des eaux) classique. À partir de ces données *discrétisées*, des techniques ont été mises au point pour caractériser les contacts grain-à-grain mais aussi pour mesurer la cinématique de tous les grains identifiés entre les états pour lesquels des images tridimensionnelles existent. La cinématique des grains est mesurée avec deux outils spécialement développés: "ID-Track" suit les grains et donne leurs déplacements; cette information est ensuite nécessaire pour une technique hybride de corrélation d'images discrète pour mesurer la rotation de chaque grain.

Des mesures à l'échelle du grain sont présentées en détail pour un essai, et sont ensuite comparées à des essais dans des conditions différentes. L'objectif est de mettre en évidence les micro-mécanismes responsables des différents comportements macroscopiques observés pour ces matériaux. Cette comparaison met en évidence certains micro-mécanismes importants tels que, par exemple, la déformation d'un échantillon. Celle-ci est concentrée dans une bande de cisaillement bien développée qui correspond à une bande de grains avec des rotations intenses. Dans un échantillon de grains arrondis, cette bande – définie par des grains – est très nette. Par contre, avec des grains anguleux, les rotations des grains dans la bande "polluent" les grains voisins; leur forme facilite l'engrenage entre grains, qui limite aussi les grosses rotations individuelles des grains. Cette différence de mécanisme de déformation est utilisée pour expliquer la différence dans la contrainte résiduelle observée à macro-échelle. Des signes de déformation localisée sont systématiquement observés en train de se produire avant le pic de déformation des échantillons et des systèmes complexes de chaînes de rotations de grains (qui correspondent aussi à une mesure locale de déformation calculée avec les déplacements des grains) sont remarquées autour du pic de la réponse macroscopique de chaque échantillon.

Contents

1	Introduction	12
1.1	Continuum models for soil behaviour	12
1.2	Localisation in granular materials	12
1.3	Micro-mechanical information from Discrete Element Models	14
1.4	Experimental micro-mechanics in 2D	15
1.5	Experimental micro-mechanical observations in 3D	16
1.6	Structure of this thesis	16
2	X-rays for imaging geomaterials in three dimensions	20
2.1	Basics of x-rays for imaging of geomaterials	20
2.1.1	X-ray radiography: historical perspective	21
2.1.2	Tomography	21
2.1.3	Basic properties of x-rays	22
2.2	X-ray scanner in Laboratoire 3SR	23
2.2.1	Technical details of x-ray source and image acquisition system	25
2.2.2	Tomographic scans	28
2.2.3	Reconstruction of a 3D field of x-ray attenuation	29
2.3	Differential displacements in the 3SR acquisition system	33
2.3.1	Preliminary study	33
2.3.2	TomoFrame [®]	35
2.4	Conclusions	40
3	Experimental campaign	41
3.1	The triaxial compression test	42
3.1.1	Specimen size	43
3.2	Materials tested	43
3.2.1	Hostun sand	43
3.2.2	Ottawa sand	44
3.2.3	Caicos ooid	44
3.2.4	Conclusions	46
3.3	Specimen preparation	47
3.3.1	Preparation of membrane for pluviation	47
3.3.2	Deposition of grains into membrane	49
3.3.3	Removal of mould	50
3.4	μ Tomo triaxial apparatus	51
3.4.1	Loading system	51
3.4.2	Triaxial cell	52
3.5	Testing procedure	54
3.5.1	Test setup	54
3.5.2	Triaxial testing	56
3.5.3	Naming convention	56

3.6	Experimental campaign	57
3.6.1	Hostun sand	58
3.6.2	Ottawa sand	62
3.6.3	Caicos ooids	65
3.7	Representativity of specimens and their responses	69
3.8	Summary of mechanical responses	70
4	Image analysis	72
4.1	Processing of individual 3D images for the identification of grains and contacts	73
4.1.1	Binarisation	74
4.1.2	Separation (or segmentation)	75
4.1.3	Labelling	81
4.1.4	Definition of contacts	83
4.2	Measurement of grain and contact properties	85
4.2.1	Measurement of particle positions	86
4.2.2	Measurement of particle orientation	94
4.2.3	Measurement of particle volume and surface area	95
4.2.4	Measurement of particle width and length	96
4.2.5	Characterisation of contacts	96
4.2.6	Saving measurement information	101
4.3	Porosity (and void ratio) measurements	103
4.3.1	“Continuum” porosity measurement	104
4.3.2	Grain-based porosity measurements	106
4.4	Measurement of specimen volume	107
4.5	Conclusions	108
5	Following processes in three dimensions	110
5.1	Principles of Digital Image Correlation (DIC)	111
5.1.1	Cross correlation	112
5.1.2	Refinements for subpixel and rotational degrees of freedom	113
5.1.3	Summary	113
5.2	“Continuum” kinematic measurements	114
5.2.1	Implementation	114
5.2.2	Strain calculation	115
5.2.3	Output of data	115
5.2.4	Summary	115
5.3	Measurement of grain kinematics	116
5.3.1	CMV-3D	116
5.3.2	ID-Track: introduction	117
5.3.3	ID-Track: measurement of displacements	123
5.3.4	ID-Track: measurement of rotations	123
5.3.5	ID-Track: measurement of rotations from grain orientation	124
5.3.6	ID-Track: measurement of rotations from grain-based 3D DIC for rotations	125
5.3.7	ID-Track: experimental characterisation of measurements	133
5.3.8	ID-Track: potential improvements	138
5.3.9	ID-Track: summary	141
5.3.10	Mathematical tools for 3D rotations	141
5.4	Contact kinematics	143
5.4.1	Algorithm for calculation of contact kinematics	143
5.4.2	Analysis of contact kinematics	145
5.4.3	Conclusion	146
5.5	Micro-strain from grain kinematics	146

5.6	Conclusions	146
6	Results and discussion	148
6.1	Experiment COEA01 ($\sigma_3 = 100$ kPa)	149
6.2	COEA01: analysis of single configurations	150
6.2.1	Global measurements: specimen-scale	150
6.2.2	Mesoscopic measurements: local porosity	151
6.2.3	Evaluation of initial inhomogeneities	154
6.2.4	Micro-scale measurements: grain sizes	156
6.2.5	Micro-scale measurements: grain orientations	157
6.2.6	Micro-scale measurements: evolution of contacts	158
6.3	COEA01: following processes in 3D	162
6.3.1	Grain displacements	162
6.3.2	Grain rotations	167
6.3.3	Axes of rotation	168
6.3.4	Calculation of local strain	169
6.4	COEA01: discussion of processes observed	172
6.5	COEA01: conclusions	174
6.6	Comparative analysis of tests on three different materials	174
6.6.1	Discussion of macroscopic responses	174
6.6.2	Comparison of tests COEA01 and COEA04	177
6.6.3	Residual stress states	181
6.6.4	Development of shear bands	189
6.6.5	Initial Kinematics	195
6.6.6	Conclusions of analysis of other tests	197
7	Example applications	199
7.1	Preparation for grain-based discrete elements	199
7.2	Application of graph theory to analyse grain kinematics	202
7.2.1	Community structures	203
7.2.2	Closeness centrality	206
7.3	Conclusions	207
8	Summary, Perspectives and Conclusions	209
8.1	Summary	209
8.2	Perspectives	212
8.2.1	Improvement of tools for micro-mechanical characterisation	212
8.2.2	Further investigations with the developed tools	213
8.2.3	Investigation of micro-scale mechanisms requiring further developments	214
8.2.4	Future work	216
8.3	Conclusions	217
9	Résumé en Français	218
9.1	Introduction	218
9.2	Analyse d'images	219
9.3	Quelques résultats importants	223
9.4	Conclusions	229

Nomenclature

Definitions of some Key Words

A **grain** will be used (naturally) to specifically denote grains of sand when talking about geomaterials.

A **particle**, as opposed to grain will be used to denote any uniquely defined digital 3D object. This means that in a given 3D volume each particle must be distinguishable from other particles, and from the background. This is often achieved by assigning a unique, positive integer number to each particle (its label), and setting all voxels belonging with a different number. It would in theory also be possible to maintain a list of particles external to the 3D volume.

A **voxel** is a 3D pixel. When the pixel size of a given scan or reconstructed 3D volume is given as, for example, $15\mu\text{m}$ per pixel, this means that the voxels measure $15 \times 15 \times 15 \mu\text{m}$. In this work, voxel will be used to describe the 3D equivalent of a pixel as seen above, and pixel will be used in an undistinguished manner for 2D and 1D pixels – rationally a measurement of distance should be given in linear pixels (lixels?).

A **radiogram** is a 2D image of x-ray attenuation, made with an x-ray source and recorded on a detector or imager. Since x-rays travel to some extent through most materials, a radiogram is a depth-averaged image over the path of the x-rays to the detector. See Figure 1.

A **tomograph** denotes an x-ray scanner which allows imaging of the specimen at different angular positions. In the medical domain this is typically achieved by rotating the x-ray source and detector around a patient. In this work however the “patient’s” well-being is less of a concern, so for convenience, the source and detector remain fixed in space, and the specimen is rotated.

A **tomographic scan** or **x-ray scan** refers to a series of acquisitions of radiographs (see above) of a given object, from different angles. These are typically acquired with the objective of “reconstructing” these into a 3D image - see “reconstruction” below.

Full-field measurements are those that evaluate a field of a given quantity all over the area



Figure 1: A radiogram (a 2D image as defined above) of a human leg. Bone attenuates x-rays more strongly than the surrounding flesh, and appears lighter in this inverted image, which is standard

of interest, as opposed to measurements on the boundary of the area of interest. These kind of measurements are typically, but not necessarily, non-destructive (FIB and BIB are examples of destructive 3D imaging techniques).

Reconstruction or **tomographic reconstruction** refers to the reconstruction of a 3D model of x-ray attenuation coefficient reconstructed from a tomographic scan (see above).

“3D images” is used to denote all full-field information (see directly above) measured over considered in this work as representing the contact a 3D volume composed of **voxels** (defined above). Most of the 3D images used in this work derive from reconstructed tomographic scans. These 3D images, are of course just large 3D matrices, where each element contains some information, such as a measured full-field property (typically “reconstructed x-ray attenuation coefficient” in the images that serve as the input to all the image treatment described here).

Greyscale is used in the computer science sense to denote a data type – greyscale images, pixels or voxels are ones which can carry a range of different values (greylevels), often in this work of reconstructed x-ray attenuation coefficient.

Binarisation is the process used in order to turn a greyscale image into a binary image – one in which there are only two values (usually grain and void in this work). This is often achieved by the use of a threshold grey value, above and below which the two phases are assigned.

Segmentation literally means splitting apart, but in the world of image treatment, it is used in two related but different ways. It is sometimes used to mean “binarise” (where the different phases present in the image are split apart from each other). Here segmentation is used in a more physical sense to mean the separation of objects *from each other* in the image. To this end a “watershed”-type algorithm is generally used. For clarity separation will be used where possible.

In-Situ is used, for this work to mean experiments performed wholly inside an x-ray scanner (*i.e. with loading performed inside the scanner*). This is opposed to the standard geotechnical meaning of *in-situ* measurements: which are (typically mechanical) tests performed in the field, as opposed to in the lab.

Description of Acronyms and Abbreviations

TXC - Triaxial Compression Test, described in the chapter detailing the experiments.

DIC or DVC - Digital Image (or Volume) Correlation, see Section 5.1 for an introduction.

ID-Track - name given to a technique developed in this work and published in Andò *et al.* (2012a) and Andò *et al.* (2012b)

NCC - Normalised Correlation Coefficient, a mathematical tool to evaluate the similarities between two signals, or matrices.

Chapter 1

Introduction

This chapter introduces the subject area of this doctoral work, putting it into historical perspective as well as briefly describing recent work in the most closely fields connected. After having laid down the historical context the structure of the work itself is then introduced in Section 1.6.

1.1 Continuum models for soil behaviour

The response of soils to mechanical loading has been of interest to civil engineers ever since substantial foundations have had to be built or large loads have had to be applied to soil. The practical information desired by engineers is typically predictive *e.g.*, the amount of settlement that will happen or whether a slope or wall will fail in a given loading scenario.

A fundamental step required in engineering-scale soil mechanics problems is to abstract forces and displacements to stresses and strains applied on a continuum. Continuum models for soil behaviour which are able to capture hardening plasticity have had a lot of success in reproducing the mechanical responses of some soils. The modified CamClay model developed in Cambridge by Roscoe and Burland (1968) is a particularly good example of this.

Continuum models of soil behaviour are ideal for use in engineering-scale Finite Element simulations. Finite Element simulations allow prediction of the effects of the complicated geometries and stress paths in a given site. However, the power of this kind of modelling depends critically on the quality of the model used to describe the soil's behaviour.

In order to inform any kind of model, experiments must be carried out on the soils that need to be modelled. Since soils are pressure-dependant materials (unlike metals), they need to be tested under pressure, which in laboratory testing is typically applied by means of a pressure cell – which limits access to a specimen once it is set up inside the cell. The parameters required in classical continuum models relate to the stresses and strains measured in a specimen which is considered as a representative elementary volume, or material point of a continuum (*i.e.*, its stress state can be described by a single tensor, instead of a field of a tensor quantity); subsequently classical soil mechanics testing has limited itself to measurement on the boundaries of a specimen. This in turn has historically limited observations of deformation that occurs *within* a given specimen, since the averaged specimen response is the only behaviour that can be taken into account in conventional continuum models.

1.2 Localisation in granular materials

As elegantly stated in the opening of Desrues (1984): “*Strain localisation is a phenomenon that is both commonplace as well as fascinating*” (translated from French). The phenomenon is described as commonplace due to its frequent apparition in the study of the mechanics of geomaterials, as well as in other domains dealing with granular materials. The fascination of strain localisation lies in the fact that materials that start deforming relatively homogeneously are able to transition

into a highly inhomogeneous (localised) mode of deformation. The apparition of shear bands, the most comment form of localisation in granular materials, is strongly related to failure, and is consequently a subject of great importance, not only for research but also for engineering practice.

Shear bands have been observed experimentally as zones of a specimen that have intense shear strain (as well as volumetric strain in dilatant materials like sand) and outside of which the material has much less strain. Shear bands have been observed to occur in planar regions with a small, but non-zero thickness. The angles of these planes have been objects of a lot of study (see Vardoulakis and Sulem, 1995). In sands, shear bands have been observed to be in the order of 10 to 20 grains thick, see Figure 1.1 from Oda *et al.* (2004), for an example.




Figure 1.1: Post-mortem observation of shear band thickness in a sand specimen deformed under triaxial compression (from: Oda *et al.*, 2004, Original caption: Microstructure in a shear band: (a) magnified image (from x-ray tomography); (b) sketch to emphasise the column-like structure in a shear band)

Strain localisation in soils causes difficulties for continuum models. Terzaghi (1920) states: “...Coulomb... purposely ignored the fact that sand consists of individual grains. Coulomb’s idea proved very useful as a working hypothesis... but it developed into an obstacle against further progress as soon as its hypothetical character came to be forgotten by Coulomb’s successors. The way out of the difficulty lies in dropping the old fundamental principles and starting again from the elementary fact that sand consists of individual grains”. A perfect example of the obstacles that Terzaghi mentions is the inability of classical continuum models to adequately capture failure. Strain localisation is in many (if not all cases) the mechanism of deformation that leads to failure in granular materials.

If a simple lab test such as a plane strain compression test is reproduced in a finite element simulation using a classical continuum model with strain softening, then strain will concentrate into a shear band. This means that strain localisation can be qualitatively reproduced. The shear bands that result from these simulations are, however, *mesh-dependent*. This means that the thickness of the mesh used to approximate the continuum has an effect on the shear band thickness obtained (*e.g.*, Bažant *et al.*, 1984). Mesh dependency of shear band thicknesses also has an effect on macroscopic response; which makes the simulation not objective at either scale.

In order to resolve the problem of non-objectivity in shear band modelling, enriched continua have been developed. Enriched continua enlarge the scope of the direct relationship between stresses and strains, which is implicit in classical continua, by taking into account higher order quantities (such as the second spatial derivatives of the displacement – discussed later). The result of this is the introduction of a so-called *length-scale* to the continuum, which is related to and controls the final dimensions of the shear bands that are obtained in simulations. Different types of enrichment are possible: local higher order continua preserve the idea of “simple materials”, as defined by Truesdell and Noll (1965) as continua where the state of a given material point

depends on the history of stress and strain in that point alone. The relationship between stresses and strains is thus enriched locally by adding dependence either on the rotation of material points (Cosserat Continua, first defined by the Cosserat brothers: Cosserat and Cosserat, 1909) or on the second derivative of material points' displacements in so-called second-gradient models (see Germain, 1973). Second gradient models have been particularly studied and developed in Grenoble (see Chambon *et al.*, 2001; Matsushima *et al.*, 2002, for applications to localisation) – recent work (Chambon, 2011) provides a general overview. Non-local methods also exist, which use a definition of a material point consistent with Truesdell and Noll (1965)'s definition of non-simple materials, that include the history of neighbouring material points in the current state of stress and strain. Enriched continua can be thought of as a first step towards recognising the fact that (as the quotation of Terzaghi clearly states) a material's small-scale – generally known as a micro-scale (although not necessarily on the scale of micrometers) – is important.

Enriched continua correctly reproduce shear band widths in a mesh independent way, but of course still require constitutive models to relate stresses and strains for a given material. Phenomenological constitutive models can continue to be used and calibrated for these enriched continua. However, shear bands in sands (as shown in Figure 1.1 and confirmed in many sources) appear to have a characteristic length (their width) in the order of $10 \times D_{50}$. The modelling of such a small number of grains as a continuum seems relatively tenuous. Recent modelling efforts, starting from this observation, have endeavoured to keep the finite element framework and have focussed on replacing classical constitutive relations with simulations of a veritable micro-scale. These models are known as multi-scale models. Two-scale, or double-scale models come in two principal varieties, Finite Elements at the macro-scale with Discrete Elements at the micro scale, called FEM-DEM models (examples of which can be found in Andrade *et al.*, 2011; Nitka *et al.*, 2011) and FEM² or FE² models, where the micro-scale is modelled by a finite element simulation of the micro-structure (see, for example, Frey *et al.*, 2012; Kouznetsova *et al.*, 2001). These multi-scale approaches obviously still require enriched continua at the macro-scale if problems of mesh dependence are expected, which is the case as soon as the multi-scale approach results in a softening behaviour at the macro scale.

Another important class of continuum models that take into account the effect of micro-structure are models that use fabric (orientations of grains, contacts and voids – often expressed as a fabric tensor) as a state variable. These models do not necessarily need enriched continua. The model proposed in Fu and Dafalias (2011) is an example of this.

This section has shown the progressive introduction of approaches that recognise the importance of micro-structure in modelling of behaviour of granular materials. These important and various theoretical developments need micro-mechanical experimental observations in order to be implemented, tested and verified. These will be covered in the next section.

1.3 Micro-mechanical information from Discrete Element Models

The modelling of strain localisation – described above in Section 1.2 – has required the development of advanced tools to take into account the effect of granular material's micro-scale. These models require micro-scale observations for calibration and testing. In the simplest case, enriched continua need observations of shear band thicknesses in order for a physically relevant internal length to be set. Multi-scale models, and models that depend on a fabric tensor require much more detailed micro-mechanical observations to be made and followed as a granular medium deforms.

Such advanced modelling tools require advanced experiments to be undertaken in order to follow the micro-mechanical evolution of granular material as it deforms. These experiments differ so radically from the traditional soil mechanics experiments (typically performed to inform continuum models) that micro-mechanical experimental observations were historically lacking when these models were developed. Historically this gap has been filled by Discrete Element Models

(DEM, see O’Sullivan, 2012, for a review). Very briefly DEM allows particle-level simulations of granular materials with simple contact laws between particles. When a sufficient number of particles are simulated, interesting trends can be observed and measured. Since each particle is known to the simulation, coordination numbers, displacements, rotations and forces can easily be obtained for every particle in every step of the simulation. For reasons of computational efficiency, by far the most common shapes to simulate are circles in 2D and spheres in 3D. Some time has been dedicated to describing grains of more complex shapes, either by allowing ellipses, polygons or by glueing several spheres together (with and without overlap). These simulations have allowed many grain-scale phenomena to be revealed and studied, but they remain simulations, and comparisons to reality are always necessary. Currently comparisons to reality have only been made using macroscopic responses, which are typically hard to match.

1.4 Experimental micro-mechanics in 2D

With micro-scale experimental observations and measurements required and desired both by the discrete element community, and by the users and developers of enriched continua, experimental micro-mechanics has not been at a standstill. Some fascinating experimental equipment has been designed that allows the full-field measurement of grain or particle properties. This started initially with 2D observations – typically by time-lapse photography – of deforming media in plane strain. Plane strain apparatus can be fitted with a glass window to allow the specimen to be observed and photographed in the plane of plane strain. Photos can then be compared in order to measure a field of displacements between these two configurations. The very comprehensive paper by Desrues and Viggiani (2004) details results of an extensive experimental campaign carried out in Grenoble with measurements obtained by stereophotogrammetry – the work of many other groups is cited therein. These measurements characterise localised strain in the specimens tested, but are however not at the grain scale.

A rather unique machine developed in Grenoble, called $1\gamma 2\varepsilon$, allows 2D tests on “Schneebeli materials” (3D analogues of a 2D material, *i.e.*, 2D shapes with an extruded third dimension) with full control of the imposed strain in 2D. Its faces are open, allowing photos to be taken of the Schneebeli particles rearranging. Particle sizes can be chosen such that individual grains are easily identifiable in the images acquired. This allows 2D experimental observations of kinematics in the style of 2D DEM simulations. Work on this machine has revealed, and confirmed, for example, the importance of contact orientations in the deformation of granular materials. Sibille and Froiio (2007) pertinently state: *“The development of discrete element methods (DEMs)..., provided powerful and flexible investigation tools. On the other hand, this has had the unfortunate effect that relatively few attempts have been made to develop finely-tuned experimental techniques for microscale investigations of granular materials. This has lead to the paradox of micromechanics of granular materials as a science based almost entirely on «virtual evidence»”.*

Calvetti *et al.* (1997) lay out challenges for the experimental characterisation of micromechanical behaviour as follows:

“The full description of micromechanical behaviour may be divided into three points:

- 1. Description of the structure, that is to say, position of grains and contacts between them.*
- 2. Description of the kinematics evolution: displacements, rotations, evolution of contacts.*
- 3. Description of intergranular forces.”*

Inference of forces in 2D is possible with the use of photoelastic disks, for example. Photoelastic materials are birefringent when strained. This can be imaged with polarised light, which in turn means that strain can be measured at a particle level, and stress (and therefore forces) can be evaluated using a constitutive hypothesis. This allows phenomena such as force chains between particles to be seen in real, plain strain, experiments – see for example Majmudar and Behringer (2005) and Lesniewska and Wood (2011).

1.5 Experimental micro-mechanical observations in 3D

The passage to full field measurement of 3D granular materials adds a considerable technical challenge. Since micro-scale measurements are required as a material evolves, destructive image techniques that progressively acquire 2D images as an object is cut away, are simply not applicable to the *imaging of a process*.

This means that *non-destructive* full-field 3D imaging is required in order to be able to achieve the goal of imaging a process. X-ray tomography, a technique principally developed in the world of medicine (where destructive imaging of live patients is not really an option), is just such a technique. X-ray tomography is extremely well adapted to material science and is currently booming in material science research.

X-rays are able to travel through matter, and are increasingly absorbed by materials with higher atomic numbers – which are vaguely correlated to density. Acquiring many 2D x-ray transmission radiographs while rotating an object allows a 3D field of attenuation to be reconstructed inside the scanned domain. In the pioneering work of Desrues *et al.* (1996), x-ray tomography is used to scan full-size (100 mm diameter) triaxial specimens of sand during deformation. Due to the limited spatial resolution of the device used, individual grains are not visible. The results of this work show that even though shear band might not be visible on the outside of a specimen, when one is able to look inside, complex patterns of localised deformation can be seen thanks to their characteristic density.

Since the pioneering work of Desrues *et al.* (1996), x-ray tomography has advanced considerably; technological developments specifically for materials science have traded off increased scanning times (since x-ray dosage is not a problem for most materials) against a great gain in spatial resolution. Furthermore, synchrotrons, such as the ESRF in Grenoble, are also sources of extremely powerful and high-quality x-ray light, which also provide increasingly fine spatial resolution at considerably reduced scanning times. These gains in spatial resolution have made x-rays an ideal tool for making observations at the grain scale for coarse-grained granular materials. The sensitivity of x-rays to density means that grains and pores can be easily distinguished, if the spatial resolution of measurement is sufficient. This can be seen in the x-ray tomography images in Figure 1.1, left (an example from this work can be seen in Figure 4.1).

Many groups have started to use x-ray tomography to investigate particulate matter, the success of the recent cycle of international workshops on X-ray Computed Tomography for Geomaterials (GeoX) conferences (see Alshibli and Reed, 2010; Desrues *et al.*, 2006; Otani and Obara, 2003) goes a long way to prove this.

X-ray imaging at the grain scale during deformation has been performed by other groups, including notable examples from Matsushima *et al.* (2006) and Alshibli and Alramahi (2006). These have proved the interest of grain-scale *in-situ* imaging, and have paved the way for the principal inspiration for this doctoral work which is the test on Hostun sand published and analysed in Hall *et al.* (2010). This test has been scanned *in-situ*, giving 3D images of the specimen deforming, where individual grains are clearly visible in the acquired 3D images. These time- or strain-lapse 3D images are locally analysed both as images of a deforming continuum, as well as at the grain-scale, to provide some unprecedented measurements.

1.6 Structure of this thesis

This thesis' primary objectives are the development of tools that allow micro-scale characterisation of granular media as they evolve in 3D. In this work x-ray tomography is systematically applied to the imaging of small specimens of different granular media, as they are deformed in a specifically designed triaxial cell. The multiple 3D images coming from each single test then need to be treated in order to make measurements of the specimen's deformation mechanisms at different scales. The idea is to cover and extend the bases set out in Hall *et al.* (2010) and to

more systematically study a range of different materials in different configurations.

Chapter 2: X-rays for imaging geomaterials in three dimensions

Chapter 2 provides a brief introduction to the most important parts of the physics of x-rays, as they apply to this work. X-ray tomography – the technique whereby many x-ray radiographs of an object are acquired at different angular positions, allowing its x-ray attenuation field to be reconstructed in 3D – is then briefly introduced.

The Laboratoire 3SR x-ray tomograph, used in this work, is described in Section 2.2. Technical details of the source and acquisition system, as well as the method by which the 3D images are reconstructed from the acquired radiographs of the specimen rotating are provided.

A simple tool to measure the quality of reconstructed images was implemented, thus permitting the correction of the acquisition geometry during reconstruction and, in turn, allowing an improved image quality to be achieved. The application of this correction to a progressive series of scans during the same test reveals that the geometry of the acquisition system is not stable with time, since the correction applied varies between progressive reconstructions of the same test. Some hypotheses about the source of these instabilities are made, and tested with a preliminary test over many hours. This preliminary test clearly shows a correlation between measured differential displacements and long term variations in temperature in the room. A tool (called TomoFrame) was then specifically developed with the double objective of allowing a more precise study of thermal effects, as well correction of radiographs during particularly long tomography scans. An example application of the measurements and corrections possible is given at the end of the chapter.

Chapter 3: Experimental campaign

Chapter 3 details the experimental part of this doctoral work, and starts by giving a brief general introduction to the triaxial test in section 3.1. The three different materials studied in this work (Hostun sand, Ottawa sand and Caicos ooids) are then detailed in Section 3.2, giving where possible grain size distributions and thus D_{50} values for the materials as well as very high resolution images of a few grains in order to qualitatively assess their morphology, which is increasingly rounded in the order in which they are listed.

The specimens tested measure 11 mm diameter and approximately 22 mm height – this extremely reduced size is chosen so that the enlarging effect of the conical x-ray beam can be used to zoom the image of the grains to such an extent that individual grains are visible in the reconstructed 3D images coming from tomography. The technique adopted for specimen preparation (detailed in Section 3.3) is based on standard preparation techniques, but has had to be adapted and modified due to the small size of the specimens desired. The preparation of such small specimens is of considerable technical difficulty. Specimens are tested dry, are only deposited by pluviation and loading platens are not lubricated. Thereafter the triaxial cell used is detailed in Section 3.4. The triaxial cell, which is detailed in Section 3.4, also differs from normal cells, principally due to its size, but also since it is entirely made of PMMA (which is relatively transparent to x-rays) and has no return-bars for the axial force. The procedure for the triaxial testing of specimens with *in-situ* x-ray tomography is detailed in Section 3.5.

The experimental campaign undertaken is detailed in Section 3.6, with each material tested in triaxial compression at 100 and 300 kPa confinement. Backup tests exist for all configurations except Caicos ooids at 300 kPa, where a test at 400 kPa from previous work is used to confirm the macroscopic results.

The macroscopic specimen responses, coming from a combination of global measurements from the loading system (force and displacement) and 3D image treatment (calculation of specimen volume to give volumetric strain) are shown. These responses are commented, and finally, the representativity of the specimens is discussed in Section 3.7.

Chapter 4: Image analysis

Chapter 4 details the way in which 3D images of a granular material coming from tomography are manipulated to extract micro-mechanically relevant information. Section 4.1 details how grains are defined in these images. The first key step is to reduce the image into a binary map of the solid and pore phases in the specimen domain. This is done against a physically measured quantity (grain volume, coming from grain mass). This binary map is then separated into individual grains by a separation algorithm. Once grains are separated, they are individually numbered (labelled). The process of separation also allows the convenient definition of contacts between grains.

Once a map of labelled grains exists individual grains can be extracted from the image and measured – as detailed in Section 4.2. The different measurements that can be made to characterise discrete objects like grains or contacts are then studied.

To conclude the chapter, mesoscopic measurements, such as the measurement of porosity (Section 4.3) are detailed, as well as the macroscopic measurements of specimen volume (Section 4.4).

Chapter 5: Following processes in three dimensions

Chapter 5 describes the local measurements of kinematics that can be made *between* 3D images. It starts off by briefly describing a general tool for the comparison of images: Digital Image Correlation (Section 5.1).

The continuum implementation of this tool in the code TOMOWARP is outlined in Section 5.2, which shows how “continuum” local kinematic measurements can be made with such a tool. Measurements of kinematics are made at nodes that are laid out on a regular grid. Measurements made on a regular grid simplify both the calculation of strain, as well as the subsequent visualisation of results, since a scalar quantity on a regular grid of nodes can simply be considered as an image.

The method with which individual grain kinematics are obtained is then detailed in Section 5.3. A simple algorithm called ID-Track has been developed in order to follow grains between configurations, recognising grains based on some invariant characteristics, such as volume. ID-Track requires two labelled 3D images of grains to have been produced and simply endeavours to make the link between grain labels in one configuration and grain labels in a subsequent configuration.

The measurement of grain displacements can be easily obtained by comparing the position of the centre of mass of a tracked grain between configurations. The measurement of grain rotations is more challenging due to the fact that digital images of grains cannot be oriented very reliably using classical image processing techniques. A hybrid grain-tracking and discrete image correlation approach has therefore been developed to measure grain rotations, without depending on a measurement of grain orientation. The two techniques for measurement of rotations are detailed and compared on real data.

Section 5.4 covers contact kinematics, for which an algorithm has been developed. At the time of writing, contact orientations cannot be measured with a sufficient degree of confidence in order to show preliminary results about contact kinematics.

Chapter 6: Results and discussion

Chapter 6 shows the analysis made possible with the tools developed and described in Chapters 4 and 5. These tools will be used to analyse the triaxial compression test COEA01 performed at 100 kPa confining pressure on Caicos Ooids. Section 6.2 shows the sorts of results that are possible with scalars calculated on single images. The section starts by commenting results at the specimen scale and continues onto meso-scale measurements like porosity, and finally shows grain and contact based micro-scale measurements.

Section 6.3 shows the measurements that are possible when the material is being followed. Emphasis is put on discrete, grain-based kinematical data (displacements and rotations). Displacement data is then processed into strain data, by measuring displacement gradients over a set of neighbouring grains defined by a triangulation of the grain centres. This is in turn compared to the strain fields obtained by continuum image correlation with TOMOWARP.

The Chapter continues with the analysis of the other experiments that have been detailed in Chapter 3. These experiments have tested specimens with a range of grain types at 100 and 300 kPa confining pressure. The objective of this chapter is to make new micro-mechanical observations of the deformation of granular media, and in particular the effect of grain shape and confining pressure on the micro-mechanisms observed. For the sake of readability results of the analysis of the various tests will not be shown in the same detail as done in Chapter 6, instead the important differences found in the micro-mechanical behaviour will be shown and discussed. This discussion will focus on some key events in the loading history of the specimens tested, starting from a micro-mechanical characterisation of the fully developed shear band in Section 6.6.3, followed by the developments leading to the shear band in Section 6.6.4. Finally the initial kinematics of the samples are analysed in Section 6.6.5.

Chapter 7: Example applications

Chapter 7 shows some highlights of two selected international collaborations involving the data acquired during this doctoral work. The first, detailed in Section 7.1, describes a collaboration with Jose Andrade and Ivan Vlahinić in Caltech (U.S.A.), where the final objective is to perform Discrete Element Simulations on a collection of discrete elements which are “avatars” of the real grains tested in the experimental campaign. The first step towards a smooth definition of the grain avatars has been completed successfully, and submitted to Granular Matter – some illustrative examples of the results born of this collaboration are shown.

Section 7.2 describes a second collaboration with Antoinette Tordesillas, David M. Walker and Andrew Cramer in the University of Melbourne (Australia). This collaboration focusses on using concepts from complex systems, and, in particular, recent developments from neuroscience to look for abstract trends in the kinematical grain data produced in this doctoral work. A paper has been accepted for publication in the Philosophical Transactions of the Royal Society A, from which some illustrative results are shown.

Chapter 8: Summary, Perspectives and Conclusions

Chapter 8 wraps up the work described in this thesis, drawing out the principal developments, and observations. The opportunities for future work are discussed.

Chapter 2

X-rays for imaging geomaterials in three dimensions

As described in the introduction, the full-field measurement provided by x-rays are a revolutionary tool in experimental geomechanics. X-rays are particularly well-suited to the study of geomaterials as they are a non-destructive measurement technique, which thus allows various configurations to be imaged. Furthermore, x-rays are sensitive to atomic density (correlated to mass density), which is particularly useful for granular media as this allows the pore and solid phases to be easily distinguished.

X-ray imaging is the key tool in this doctoral work, in which full-3D imaging during mechanical tests on granular geomaterials is performed. From these 3D images, deformation at the micro-scale is followed and analysed. This analysis has been made possible by recent advances in x-ray technology, and in particular the availability of laboratory based x-ray technology (as opposed to synchrotron sources such as the ESRF in Grenoble). Laboratory x-ray technology has allowed the systematic (and therefore time-consuming) study of several granular materials under different loading conditions that would have been difficult with the limited time allocations available in synchrotron sources.

Layout of chapter

This chapter describes how 3D images of a geomaterial are obtained. The chapter starts with a brief introduction to x-rays and x-ray tomography (Section 2.1), and follows on to describe the most important aspects of the interaction of x-rays with matter, in order to show that x-rays are indeed suitable for the imaging of the materials studied in this work.

This chapter continues with a description of how x-rays are generated and measured in the setup used for the experiments that form the body of this work (Section 2.2). The way in which a 3D image is reconstructed from the acquired data is then detailed in Section 2.2.3. Specifically-developed tools to address observed instability of the acquisition system with time are outlined in Section 2.3.

2.1 Basics of x-rays for imaging of geomaterials

A brief introduction to the most important features of x-rays is provided in this section with some historical context to x-rays and x-ray tomography. The section then continues onto a description of the most important characteristics of x-rays and their interaction with matter for the imaging of geomaterials. This subject is covered in much more depth in many textbooks, a particularly complete one for the physics of x-rays is “Elements of Modern X-ray Physics” – Als-Nielsen and McMorrow (2011), and for x-ray tomography: “Computed Tomography Principles, Design,

Artifacts, and Recent Advances” – Hsieh (2009). Both these textbooks are used as a reference throughout.

2.1.1 X-ray radiography: historical perspective

X-rays are a high-frequency form of electromagnetic radiation (with a wavelength between 0.01 and 10 nanometers) that were discovered at the end of the 19th century by Wilhelm Conrad Röntgen when experimenting with vacuum tubes. The radiation coming from the vacuum tube was found to cause a fluorescent screen to light up, even with paper and wood in the way of the beam, the beam was however stopped by metal. Röntgen then tried to photograph the beam – asking his wife to put her hand between the x-ray source and a photographic plate – thus making the first-ever radiograph, reproduced in Figure 2.1 (Glasser and Boveri, 1933).



Figure 2.1: “Hand mit Ringen” – First radiograph, showing Anna Röntgen’s left hand with a ring, by Wilhelm Röntgen (from Wellcome Trust, in Public domain, via Wikimedia Commons)

The radiograph in Figure 2.1 shows details of the bones inside the hand. The bones are distinguishable in the image because they stop x-ray transmission more than the flesh that surrounds them. As a result, in the regions behind the bones, fewer x-ray photons interact with the photographic plate being used to acquire the image, causing them to appear darker. The (presumably metallic) ring has the same effect – by blocking the x-ray beam, it appears dark in the radiograph. An example of a modern medical radiograph (with reversed colours) is shown in Figure ??.

An X-ray radiograph is therefore a two-dimensional measurement of the amount of x-ray photons arriving at each point on a detector during a given exposure time, which is an integration of the x-ray attenuation of the matter being traversed along the path of the x-rays.

X-rays were recognised immediately as a very valuable tool for medical diagnosis, and it is this field that caused x-ray technology to develop rapidly. The x-ray sources currently used in medical applications (as well as in the scanner in Laboratoire 3SR) are direct descendants of the vacuum tubes that Röntgen used for the discovery of x-rays.

Due to Bohr’s wave-particle duality, it is sometimes useful to refer to individual x-ray photons, and at others it is useful to refer to a wave or a beam. Both these terms are used (albeit not interchangeably – since they refer to two different ways of thinking) throughout.

2.1.2 Tomography

For many applications, such as assessing fractures in bone, radiography is sufficient as a diagnosis tool. However, for the truly 3D medical issues, such as diagnosis of brain problems, simple radiographs are not sufficient. X-ray tomography (or CAT – [Computer, Computed, Computerised]

Aided, Axial | Tomography scanning) was developed in the 1960s in order to address this need. Very briefly, tomography is a method for reconstructing a 3D field of x-ray attenuation coefficient inside an object starting from a series of different projections (radiographs) of the object. The different projections are in general acquired by a relative rotation of object and image acquisition system. In medical applications the source and detector rotate, in industrial/laboratory applications such as this one, the specimen is usually rotated – with a vertical axis of rotation – by a rotation stage. This technique described in a more detail in Section 2.2.3 and fully detailed in a number of textbooks on tomography, such as Hsieh (2009).

Developments since the 1960s on the medical front have been pushing towards ever smaller doses of x-ray radiation for patients, along side the ability of resolving ever-smaller objects. X-rays for material science have also pushed development towards increased spatial resolutions. Material science has also benefited greatly from x-ray radiation coming from synchrotron sources, which offer a much better (more “brilliant”) source of x-rays, which is not suitable for scanning patients.

2.1.3 Basic properties of x-rays

X-rays occupy a specific range of wavelengths in the electromagnetic spectrum. Inside this range x-rays are divided into “hard” and “soft” x-rays. Hard x-rays are on the lower wavelength end of the x-ray spectrum and consequently have higher energy which provides greater more capability for penetrating matter.

When x-ray photons interact with matter, they may do so through many different mechanisms (such as Compton Scattering, Refraction and Reflection, Pair Production, Raleigh Scattering), but for the materials studied in this work, and the energy of x-ray radiation used, the overarching mechanism of interaction of x-rays with the material is photoelectric absorption.

In a photoelectric absorption event, an x-ray photon is absorbed by an atom, which causes an electron to be ejected. The atom thus becomes positively charged – an ion – which is why x-rays are considered as ionising radiation. The photoelectric effect is considerably augmented in a given material if the energy of the incoming x-ray photons have at least the same energy as the electrons in the “K-shell” (the shell closest to the nucleus where electrons have the highest energy). This effect can increase x-ray photon absorption by up to ten times and is known as the K-edge effect (Als-Nielsen and McMorrow, 2011). For elements with a higher atomic number, the energy of electrons in the K-shell is higher, which explains why, for a given x-ray energy (*e.g.*, 50 keV), there is a sudden discontinuity in the attenuation coefficient of the material (see Figure 2.2). The K-edge energy of Europium ($Z=63$) is 48.52 keV, enough for the K-edge effect to work, whereas Gadolinium’s ($Z=64$) is 50.24 keV, insufficient for the K-edge effect; which explains the reduction in attenuation between these materials at this energy.

Attenuation in a given material is given by a coefficient of linear absorption μ . The intensity I (photons per area per unit time) of a monochromatic beam as it travels through a medium to a depth z is given by the Beer-Lambert law:

$$I(z) = I_0 e^{-\mu z} \tag{2.1}$$

where I_0 is the Intensity of the beam before entering the material

Furthermore, Als-Nielsen and McMorrow (2011) state that the photoelectric absorption cross-section, σ_{pe} , is related to μ as shown in equation 2.2, and that it varies with the atomic number Z of the absorbing material “approximately as Z^4 ”. It is this variation that gives the contrast between elements that makes x-rays so useful for imaging materials.

$$\mu = \left(\frac{\rho_m N_A}{M}\right) \sigma_{pe} \quad (2.2)$$

where: ρ_m is the mass density, N_A is Avogadro's number and M is the molar mass

The variation of attenuation with atomic number is illustrated in Figure 2.2, which shows the mass attenuation coefficient for x-rays of three different energies (50,100,150 keV) for various elements sorted by their atomic number, Z , coming from the NIST database (Hubbell and Seltzer, 1996). The mass attenuation coefficient shown here is like the photoelectric absorption cross-section (σ_{pe}), but takes into account all forms of material interaction (and is thus labelled σ_{tot}). It can be seen from the line in red that there is a positive correlation between atomic number and σ_{tot} , but with a strong influence of the K-edge effect at different energies.

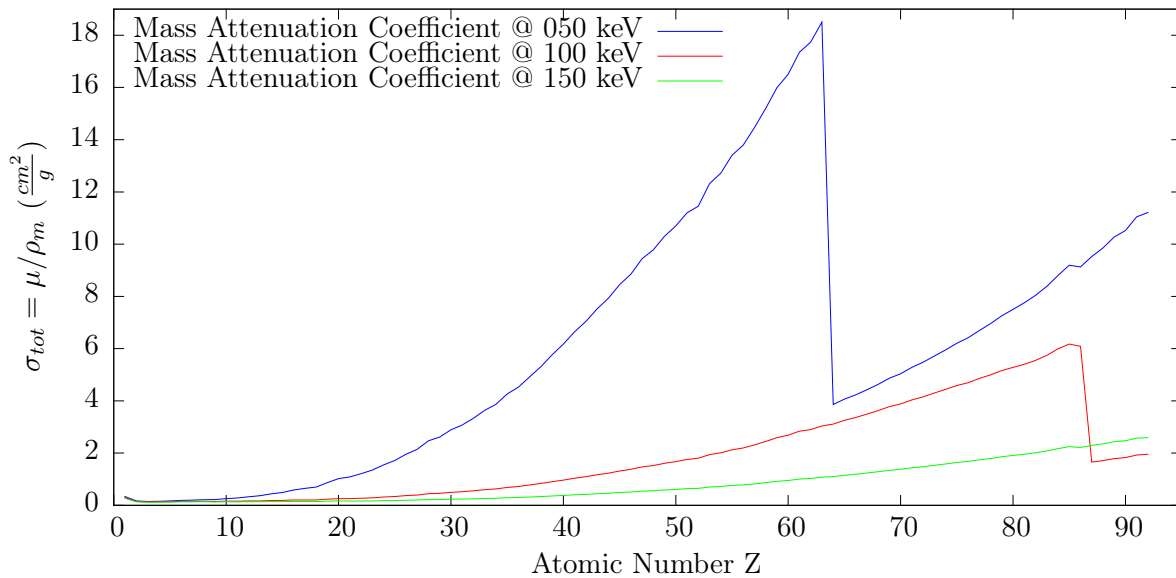


Figure 2.2: Graph showing both the mass attenuation coefficient to 50, 100 and 150 keV x-ray photons for a variety of elements. Data from the very useful Hubbell and Seltzer (1996) – available online

The different x-ray energies shown in Figure 2.2 indicate that harder x-rays have a better ability to penetrate matter than softer ones, since their corresponding mass attenuation coefficient is generally lower.

In summary, x-rays thus are a type of radiation which is able to be transmitted through matter, and which is sensitive to the atomic number of the material that it is radiating through.

2.2 X-ray scanner in Laboratoire 3SR

The x-ray scanner in Laboratoire 3SR (Grenoble) is a built-to-specification laboratory x-ray tomograph supplied by RX-Solutions (Annecy), which was bought in 2007 with funds from ANR Project MicroModEx as well as Laboratoire 3SR. It is housed inside a lead-lined cabin.

The x-ray source in this tomograph – which will be described in more detail in Section 2.2.1 – emits a cone-beam (not unlike the light coming out of a projector) in the direction of the detector, which means that as one translates a specimen from the detector towards the source, the field of view decreases, and the specimen is enlarged on the detector due to geometric magnification.

This is illustrated in Figure 2.6 and is conventionally known as zoom. This translation in the axis of the beam is provided by a trolley, running on rails attached to the wall of the cabin facing the door, visible in Figure 2.3.




Figure 2.3: Labelled photograph of the Laboratoire 3SR x-ray scanner, with background faded out for clarity. The blue edges visible left and right are the door frame of the cabin

The Laboratoire 3SR x-ray tomograph differs from standard “off the shelf” ones on a few key points:

- **Translation trolley**

The trolley, on which the rotation stage is able to translate in the direction of the beam, allows easy control of the zoom level for the specimen being scanned. Furthermore, this is a solid steel component able to carry considerable load (*i.e.*, 75 kg), which means that massive experimental equipment can be mounted onto the trolley and kept in place during an x-ray scan (provided it does not block the beam).

- **Room for experiments**

In the two directions normal to the axis of the beam (the width and height of the cabin) there is considerably more space than in an ordinary x-ray scanner. This means that large amounts of experimental equipment can be mounted on or around a specimen being scanned. For mechanical tests, for example, there is the benefit of being able to keep many data acquisition elements in the cabin itself.

- **Connection to outside**

There are two sets of chicanes, to allow the passage of tubes and cables in and out of the cabin, without risk of exposing the outside of the cabin to x-ray radiation.

- **Hollow rotation stage**

The rotation stage, which is mounted on the trolley, is a hollow ring that rotates, which allows experimental equipment to be mounted onto and, more importantly, through the ring. This means that experimental equipment can be developed in the space above and below the specimen being scanned.

- **Translatable source and detector**

The source and detector can be moved synchronously contained in the plane of the detector (*i.e.*, they can be translated in the two directions normal to the axis of the beam), which allows experimental installations of different heights to be conceived and scanned.

These features mean that the scanner is a flexible tool for the investigation of material behaviour, and specifically for this work, means that it has been possible to design a triaxial cell specifically for *in-situ* x-ray scanning (*i.e.*, scanning during a triaxial test) – see Section 3.4.

2.2.1 Technical details of x-ray source and image acquisition system

This section details the x-ray source, and the image acquisition system of the x-ray scanner in Laboratoire 3SR.

X-ray Tubes

The source used is an x-ray source with a sealed, solid-anode, micro-focus x-ray vacuum tube. A schematic of this kind of tube is shown in Figure 2.4.




Figure 2.4: Schematic showing the kind of x-ray vacuum tube (sealed, solid-anode, micro-focus) found in the source in the Laboratoire 3SR x-ray scanner

X-rays are generated by applying a large voltage (in the kV range) across the anode and cathode of the tube. The cathode contains a filament that will produce electrons on heating.

These are attracted to the anode target (made of tungsten in this case), which is responsible for generating the x-rays. Two separate mechanisms are responsible for the generation of x-rays. The first is *bremstrahlung* (German word), the braking radiation. Braking radiation occurs when electrons pass close to an atomic nucleus, that, due to its positive charge, deflects the electron, causing an x-ray photon to be emitted. This mechanism is responsible for generating x-ray photons of all energies up to a maximum of the voltage applied to the electrons, in eV. The second mechanism responsible for generating x-rays is due to the change in orbit of electrons in the atoms that make up the target. The energies of the photons thus produced are at specific energies, characteristic of the material. Both of these mechanisms generate x-rays at different energies, so the beam is called “polychromatic”.

In this kind of source, both the voltage and the current applied to the x-ray tube can be controlled. An increase in voltage applied to the tube increases the spectrum of energies created, and consequently the mean photon energy increases. Higher energy (harder) photons have an increased ability to penetrate matter. An increase in current supplied to the tube increases the number of electrons emitted, and therefore also the number of photons produced. The range of photon energies is however dictated by the voltage, so a change in current simply changes the flux of photons.

The x-rays generated in an x-ray tube leave the tube in a divergent manner, making a cone of x-ray light. This cone of x-rays in reality does not start from an infinitesimal point, but rather comes from a “spot” of a certain size. The size of the spot is typically related to the electrical power supplied to the tube. In this kind of source, the spot size is small relative to that used for medical scanning – which justifies the name “micro-focus”. The spot size is important because it must be smaller than the smallest object which is to be resolved in the radiograph acquired.

Since the x-rays generated are polychromatic, an effect known as “beam hardening” can also occur as x-rays travel through a specimen. Beam hardening literally means that the beam becomes *proportionally* harder as it crosses the specimen: because softer (*i.e.*, lower energy) x-rays are attenuated more easily by matter - as is shown by Figure 2.2. It is worth keeping in mind that (approximately speaking) the detector used (detailed below) does not distinguish between soft and hard x-rays: it measures only the intensity (number of x-ray photons per area and per unit time) of the beam, and does not measure the energy of the incident x-ray photons. If a cylinder, whose axis is in the plane of the detector, is imaged with a monochromatic (single-energy) x-ray source, the Beer-Lambert law (equation 2.1) suggests that the intensity of the beam behind the thickest part of the cylinder will be lower than at the edges of the cylinder, because the x-ray photons will have had to travel through less material (so z will be smaller). If the same image is acquired with a polychromatic source with the maximum energy equal to the monochromatic source used in the example above, the ratio between image intensity for the inner and outer part of the cylinder will be exaggerated because the softer x-rays will be attenuated more easily (since μ is expected to be higher for softer x-rays) though the widest path through the specimen. Furthermore, the x-ray beam has lost proportionally more soft x-rays than hard ones, and so has become (proportionally) harder. The fact that the ratio between the intensities has changed on the detector means that when these data are reconstructed, the coefficient of linear attenuation appears artificially lower in the centre.

In order to reduce beam hardening, filters can be used, taking advantage of the K-edge effect to attenuate, to some extent, the soft x-rays thus physically reducing the softness of the beam. In this work beam-hardening is not a big problem, since the triaxial cell used is thought to filter a good deal of soft x-rays. Software corrections for beam-hardening on the reconstructed images also exist.

Laboratoire 3SR source

The source used in this work is a Hamamatsu Corporation L8121-03. It is flexible insofar as it can apply between 40 and 150 kV to the x-ray tube, and supply currents between 0 and 500 μA

(micro-Ampères). The source operates in three different “spot size” modes: small, medium and large. The first two modes have caps on their power, but guarantee a reduced spot size.

In the work presented here, 100 kV has been selected in order to have the best contrast between pore and grain when imaging the materials studied. The highest current allowed for this voltage setting while remaining in small spot-size mode is selected as the tube current (*i.e.*, 100 μ A). The manual for the source indicates that a spot size of 7 μ m should be achieved with these settings.

Laboratoire 3SR detector

The detector used in this work is an “indirect” flat-panel Varian PaxScan[®] 2520V, marketed for medical use (see: *PaxScan® 2520XI/V* 2010). It measures the intensity of incident x-ray photons on an array of 1920 \times 1536 pixels, which can be put either into portrait or landscape mode (portrait mode is used in this work), with each pixel measuring 0.127 \times 0.127 mm. This gives a detector size of 243.84 \times 195.072mm (approximately the area of printed text on this page). This kind of detector uses a scintillator (Caesium Iodide, CsI, in this case), to convert x-rays into visible light, whereupon they are captured by the large array of photodiodes (description in *Flat Panel X-ray Imaging* 2004).

The detector works in a similar way to a typical CCD detector: it has a dynamic range which must be respected. The x-ray source is capable of producing beams of x-ray light at a considerable range of different energies and fluxes. This is accommodated for in the detector by being able to change the acquisition rate (measured in Frames Per Second) – the reciprocal of the exposure time if there is no delay between acquisitions – at which the detector reads out. If the exposure time is too long, photodiodes may saturate, meaning that none of their dynamic range is used, giving a flat, maximum value. On the other hand, if the exposure time is very short, very little of the dynamic range will be used giving a very dark image, with a low signal to noise ratio. When acquiring an image, this Frames Per Second setting is adjusted in such a way to a use as much of the dynamic range as possible (without saturating any pixels) in order to increase the signal to noise ratio of the images acquired, while also trying to keep the exposure time small, to shorten the total scan time.

For the work carried out here, the FPS was set to 1.5 (which corresponds to an exposure time of 0.67 seconds). Furthermore in a tomographic scan, eight images of each angle are acquired and averaged in order to reduce noise.

A further improvement typically applied in digital image acquisition, is that of “flat fielding”, whereby images are corrected with dark and a flat-field image (see: Seibert *et al.*, 1998). These are respectively an image captured with no exposure time (or with no incident x-rays) designed to capture the “dark currents” of the detector, and an image of a homogeneously lit scene – or in this case an acquisition with no specimen in the x-ray beam. Making the hypothesis that the detector has a linear response to x-ray intensity, this allows a good first-order correction to the acquired images. These corrections are automatically handled and applied by the acquisition program.

For illustrative purposes, Figure 2.5 shows a single projection from a tomographic scan (which is actually the result of an averaging of 8 radiographs at that angular position). The domain of the specimen, which contains calcite grains separated by air, is more highly attenuating (and therefore darker in the projection) than the porous stones top and bottom, as well as the confining fluid (water).

Zoom and spatial resolution

It is imperative to note that the number of pixels in the detector is fixed. It is however possible, as described above, to zoom the specimen by translating it closer to the source. Since there is a fixed number of pixels in the detector, the effect of zoom is to change the effective size of a




Figure 2.5: Projection of specimen (resulting from an average of 8 radiographs) at an angular position of 0° in test COEA01 in configuration 01. Brightness and Contrast of image have been modified in order to best show specimen, consequently the cell and confining fluid (water) cannot be seen. The border of the image is due to a border of inactive pixels on the detector. This image measures 1920×1536 pixels

pixel at the level of the specimen. This is best illustrated with an example: Figure 2.6 shows 2D radiographs of a leaf, acquired at different levels of zoom. All the images have the same number of pixels - since the same detector is being used, it is the effective pixel size in the image with respect to the object which is changing between these images. It is clear that when a specimen is zoomed due to the cone beam geometry, the field of view decreases, so either small volumes can be scanned at a high spatial resolution, or large ones can be scanned at a lower spatial resolution; many different scales can thus to be investigated with this equipment.

The equipment therefore has a fixed number of pixels, but has variable spatial resolution thanks to the cone beam and the ability to zoom. In this work the grain sizes to be studied are approximately known *a-priori*, which means that if individual grains need to be identifiable in the 3D images, a certain spatial resolution must be imposed.

The sands being studied have a grain size of approximately $500 \mu\text{m}$, so a *voxel size* of $15 \mu\text{m}$ was chosen in order to have around 30 pixels across the thickest cross-section of a grain, which enabled the grains to be easily distinguished in an image. Since a pixel on the detector measures $127 \mu\text{m}$, for the zoom to be such that a pixel images an area $15 \mu\text{m}$ long, the zoom factor is 8.5. This corresponds to a zoom level, on Figure 2.6, approximately equal to that of the image in the middle of the bottom row. The consequence of this zoom level is that the field of view of the detector is reduced 8.5 times, and therefore is approximately $30.5 \times 23.0 \text{ mm}$. Specimens must be designed to fit comfortably inside this field of view in order to be imaged completely.

2.2.2 Tomographic scans

The tomographic scans performed in this work consist in acquiring projections of 1024 different angular stations as a specimen is rotated around a sub-vertical axis at regularly spaced intervals from 0 to 360° . At each angular station 8 radiographies are acquired and averaged. An entire scan takes just under 2 hours (1 hour and 56 minutes).




Figure 2.6: X-ray radiographs acquired in the Laboratoire 3SR scanner of the same leaf (approximately 24 cm long) at different levels of zoom, from minimum (top left) to maximum zoom (bottom right). The images have been manually corrected with dark and flat field images, as well as locally smoothed, and adjusted for changes in overall image intensity

2.2.3 Reconstruction of a 3D field of x-ray attenuation

As stated briefly in the introduction, tomography is a technique whereby a number of different projections (radiographs) of the same object are acquired with the objective of reconstructing a field of the x-ray attenuation inside the object. Reconstruction is achieved with the use of a mathematical tool, the Radon transform, developed by Radon (1917), which describes how a function can be *reconstructed* from an infinite number of projections of the function. In this case the function of interest is the one describing the field of the x-ray attenuation coefficients (μ) inside the object being scanned, and the projections are the radiographs acquired at different angles. The objects being scanned are three dimensional, whereas radiographs are 2D; each point in a radiograph represents an integration of all the linear x-ray attenuation coefficients encountered along the path of the x-ray (as per equation 2.1).

Starting from the Radon Transform, several different reconstruction techniques have been developed, allowing a 3D image to be reconstructed from a series of x-ray projections. The simplest is perhaps the backprojection family of reconstruction techniques which starting from the first projection, simply distribute the measured attenuation equally along the path of the x-ray. This is progressively refined as subsequent projections are included. This procedure is analytic (requiring no iteration). The backprojection of a single point on a 2D slice (taken directly from Hsieh (2009)) is shown in Figure 2.7. Another important family of reconstruction algorithms are the iterative ones, such as ART (Algebraic Reconstruction Technique), which iterate in order to find the function describing the linear attenuation coefficients inside the scanned object. These techniques are slower than simple back projection, but allow a higher quality reconstruction to be made, and can model several sources of error, in order to improve the reconstruction. Recent work in this field is very promising, several examples show that where a back-projection type algorithm needs around a thousand projections, with some assumptions (such as the number of phases in the specimen) iterative techniques can produce the same signal to noise ratio in the reconstruction with less than an order of magnitude fewer projections. Strategies such as the minimisation of total variation (TV) of the reconstructed image are particularly promising

(Herman and Davidi, 2008). This opens the door for much more rapid scanning, among many other things.




Figure 2.7: Backprojection process of a single point. (a) Backprojected image of a single projection. (b)–(i) Backprojection of views covering: (b) 0 to 22.5°; (c) 0 to 45°; (d) 0 to 67.5°; (e) 0 to 90°; (f) 0 to 112.5°; (g) 0 to 135°; (h) 0 to 157.5°; and (i) 0 to 180°. From Hsieh (2009)

Practical details of reconstruction

In this work, the reconstruction software DigiCT version 2.4.2 from Digisens is used. DigiCT uses a reconstruction algorithm based on the FDK Algorithm, after Feldkamp, Davis and Kress who describe it as a “convolution-backprojection” algorithm, and provide a formulation for cone beams in Feldkamp *et al.* (1984). The principal interest of this software is that it has been written in such a way to take advantage of GPU processing (Graphics Processing Units – whose technological development has been pushed by the video games industry). This means that the reconstruction time is reduced significantly – once the correct parameters for the reconstruction are selected a reconstruction of a $1500 \times 1500 \times 2000$ volume takes under ten minutes when equipped with two NVIDIA Tesla C1060 cards. Since the reconstruction is split between two separate graphics cards, it is important that the reconstructed x-ray attenuation data are not rescaled or normalised, since this would occur on a per-card basis, creating a volume with two sets of differently normalised values. To avoid this problem, the reconstructed slices are written out in their raw, as-calculated, form. This takes the form of a 32-bit floating point RAW file.

A successful reconstruction depends on the correct setting of several parameters. By far the most important of these is the position of the axis of rotation of the object relative to the images. This is typically described by a centre of rotation (typically the middle line of the image), combined with a tilt angle from the vertical, which together give the axis of rotation.

The axis of rotation can be measured experimentally by performing a rapid tomography (around 100 projections) of a calibration tool. The calibration tool consists simply of a tube inside which are several highly attenuating spheres of approximately the same radius as the tube

– in this work spheres of radius 1 mm are used. A small distance separates all the spheres, so that when they are imaged (with the axis of the tube in the plane of the detector) they do not appear to touch each other. This instrument is placed vertically, but off-centre in the domain that is going to be scanned. A tomography of the instrument is performed with around 100 projections; in these projections the spheres have an elliptical trajectory because of the cone beam. A simple circle-recognition program, supplied with the x-ray scanner (as well as another version with the reconstruction software) can then follow the spheres, and knowing the size of the spheres scanned and by calculating the centre of each sphere’s trajectory the geometric setup of the scanner, including the axis of rotation and zoom level can be calculated.

This definition based on a pre-scan measurement of the geometrical setup of the x-ray acquisition system involves a fundamental hypothesis that it does not change after measurement. It is due to this hypothesis that as much experimental equipment is put in place before making the measurement: if the weight on the trolley that holds the rotation stage changes (see Figure 2.3), it is quite possible that the centre of rotation or the tilt will change. Experience shows that the tilt angle measured in this way is often satisfactory and can be used directly. The centre of rotation, however, in this setup, is much more sensitive to variations which can affect the reconstructed image quality. The centre of rotation is thought to drift due to small changes in ambient temperature in the cabin, this is studied in more detail in Section 2.3. In any case, the centre of rotation in the image has been noticed to drift, which means that reconstructing with the centre of rotation measured before the test is not necessarily correct for a given scan. Furthermore, and much more serious, the effective centre of rotation can also drift *during* an acquisition. Thankfully, at the zoom levels of these scans, except for rare, extreme cases (which require manual shifting of the radiographs to compensate the displacement) optimising the centre of rotation to a value which represents a mean value for the scan gives good results. A tool, suggested by Dr. Pierre Bésuelle (Chargé de Recherche CNRS at Laboratoire 3SR) was implemented in order to find the best centre of rotation based on slices reconstructed with a range of centres of rotation. Figures 2.8 middle, a) and c) show what slices reconstructed with an incorrect centre of rotation look like. An immediate observation is that the edges are defined much less sharply, when compared to the expected slice b). The idea is simply that the sharpness of boundaries (many sharp boundaries between grain and void are expected in images of dry granular media) can be quantified by measuring the amount of high frequency content in the Fourier domain of the image. The image for which the high-frequency content is the highest is likely to be the sharpest one. Figure 2.8 (top) has been calculated with a step size of 1 μm of the centre of rotation, and this allows the peak to be calculated with sufficient resolution, the peak could even be calculated more precisely by fitting a curve around the peak and finding its point of inflexion. Figure 2.8 middle b) is the slice reconstructed with the centre of rotation that gives the most high-frequency content.

Figure 2.8 (bottom) shows the evolution of the best centre-of-rotation for each scan during a given test, which shows quite conclusively that there is drift associated with time.

Optimising the centre of rotation means that the effect of the centre of rotation drift is minimised, but not eliminated. For the both the study and the elimination of these effects, a project is ongoing with Dr. Simon Salager (Maître de conférence at Université Joseph Fourier and at Laboratoire 3SR), which is described in Section 2.3.

Another key parameter that needs to be set correctly in order to have usable reconstructed images, is the beam hardening correction. Without this correction, an unequal repartition of linear attenuation coefficients is made in the reconstructed slices, with the centre of the specimen having lower values (and therefore appearing darker) than the outside. This is corrected in software, by generating a set of slices with a range of different beam-hardening correction factors, until an equal profile of grain-density is achieved for grains all the way through the specimen. If the software correction were to have proved insufficient, a hardware filter would have been added to the source.

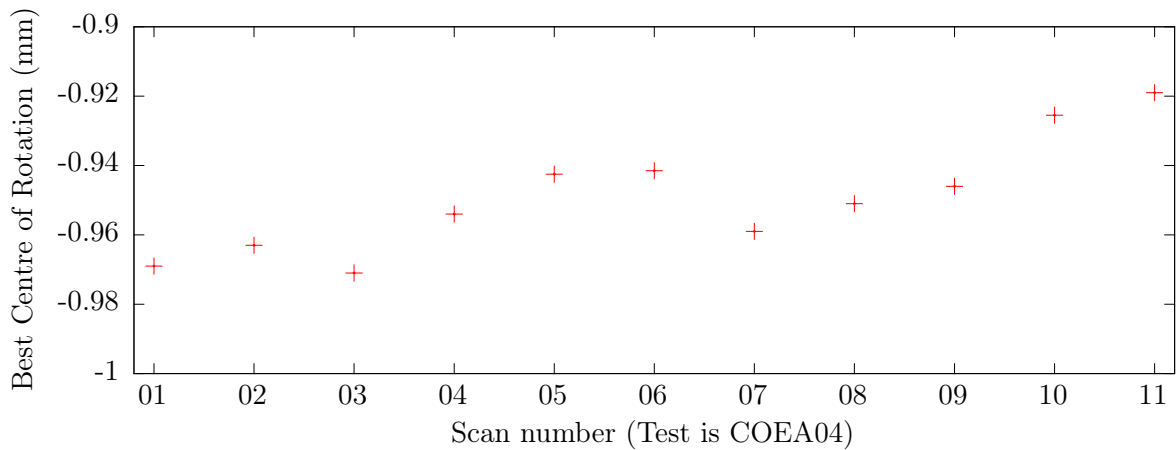

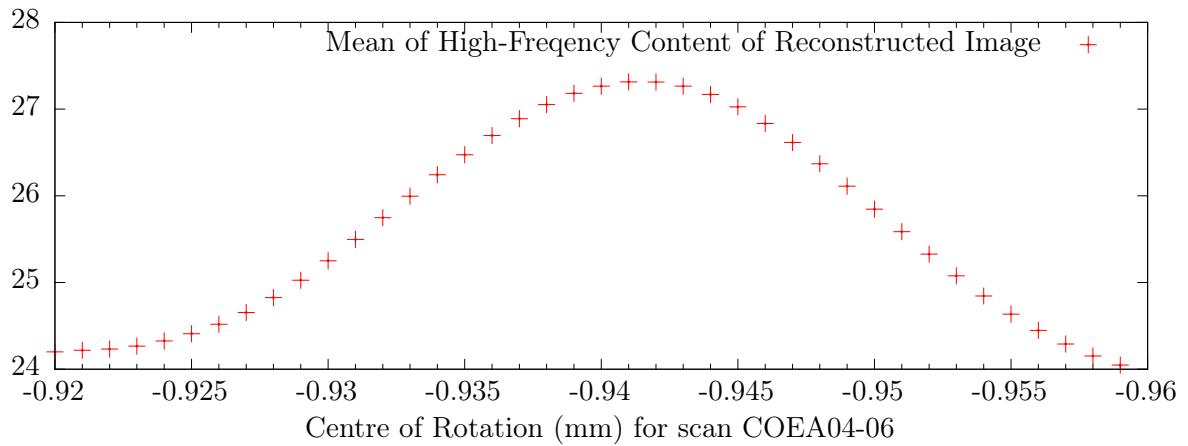


Figure 2.8: Top: Graph showing the change in the mean of the high-frequency content of the Fourier transform of slices reconstructed with a range of different centres of rotation. A global maximum is reached around a centre of rotation of -0.941 mm. Note: the change in centre of rotation between points is 0.001 mm (1 μ m), much smaller than the pixel size of the scan (15.556 μ m).

Middle: Three details of the same position in slices reconstructed with centre of rotation values equal to a) the left most point on top graph, b) the peak of top graph and c) the rightmost point on top graph.

Bottom: Graph showing the best centre of rotation, chosen as the maximum of the corresponding top graph, for each scan during a test. Note: COEA04-06 has a value of -0.941 as per top graph

Once a specimen can be reconstructed satisfactorily in an initial state, a fixed volume to be reconstructed is defined in the images and fixed for the entire test. A fixed reconstruction area is imposed for all reconstructions of a given test in order to have a fixed reference point – with the hope that despite the geometric fluctuations measured, the immobile part of the experiment (the top loading platen) does not move much relative to the origin of the reconstruction. Furthermore, the fact that all reconstructed volumes are all the same size simplifies the automation of data treatment. In the initial configuration the specimen is at its maximum height (since shortening is imposed), and so the height of the volume to reconstruct can be defined tightly. The width of the reconstructed volume, is left wide since the specimen is expected to expand in diameter, due to Poisson’s ratio, dilatancy as well as shear banding.

2.3 Differential displacements in the 3SR acquisition system

A single x-ray tomography with the settings used in this work takes approximately two hours. Since several configurations are expected to be scanned during a test, the geometrical stability of the imaging system during this time is important in order to be able to make precise incremental measurements using x-ray tomography. As Section 2.2.3 (and more specifically Figure 2.8, bottom) shows, some drift in the centre of rotation occurs during a test. Furthermore, some drift of the centre of rotation is also suspected *during* a test. Figure 2.8 (middle) shows that for the images acquired in these conditions, corrections can be applied to get a satisfactory result. However when objects need to be scanned at a higher spatial resolution (*i.e.*, bigger zoom, with a pixel size under 10 μm) this requires a correspondingly small spot size, which is achieved by reducing the power of the x-ray source. This in turn means that the intensity of the beam is reduced, so the exposure time for the acquisition of radiographs need to be increased in order to use enough of the detector’s dynamic range. A longer exposure time has the direct consequence of increasing scan time – finally meaning that scans at a higher spatial resolution are doubly susceptible to differential displacements in the image acquisition system, first due to their long scanning time, and second because at a higher zoom differential displacements have a bigger effect.

Two preliminary hypotheses are made:

1. The position “spot” of the detector fluctuates with time (communicated by the manufacturer as a potential instability), causing the radiographs acquired to only appear as if they were displaced.
2. Changes in temperature are responsible for real differential displacements (due to geometry, or different coefficients of thermal expansion) in the acquisition system.

2.3.1 Preliminary study

In order to better understand the source of differential displacements, a preliminary study was conducted. The objectives of this study were to see whether differential displacements could be measured, and if so whether the source of these apparent displacements could be ascertained. The issue was initially studied by performing a long scan (over a weekend) with the very simple setup as shown in Figure 2.9. The setup is designed to falsify the hypotheses made above, by using three steel spheres placed in such a way that they can be imaged over a long time. Each steel sphere is attached to, and expected to follow, a different part of the x-ray imaging system. As shown in Figure 2.9, sphere A is attached to the x-ray source by a G-clamp and is suspended in a block of foam in order for it to be easily identified in an image; this sphere is included to test, to some extent, hypothesis 1. Sphere B is attached to a tripod whose feet are on the floor, and which goes through the hollow rotation stage. Sphere C is attached to a block of steel that is resting on the rotation stage, to test hypothesis 2. The spheres and their holders do not touch

each other. The radiograph in Figure 2.9 shows that the three spheres are easily identifiable with this setup. It is important to note that radiographs were acquired during a weekend, with no rotation of the rotation stage. Two thermocouples were also used – one inside the cabin and one outside the cabin, both sticking upwards into the air – in order to falsify hypothesis 2.




Figure 2.9: Photo (with background faded) of preliminary setup, showing tripod resting on the floor, and g-clamp attached to the source.

Inset: Photo (with background faded) showing a frontal view of the three spheres, in their holders (the x-ray source is in the background).

Resultant Radiograph: a specimen radiograph acquired using this setup

A set of 1000 radiographs was acquired with a delay of 300 seconds (5 minutes) between them, with thermocouple readings made at the same rate. Since the objective of the test is to characterise what might be a dynamic process, the radiographs are not averaged, as this could blur displacements. The spheres were identified by thresholding, and were followed by measuring their centre of mass in the different projections acquired. This kind of image processing is covered in detail in Chapter 4. Figure 2.10 shows the changes of the measured quantities during the test (change in position in the image of the spheres and change in temperature of the thermocouples). The displacements reported in pixels can roughly be converted to microns by multiplying by 5.

Figure 2.10 shows that a temperature range of about 3 °C is witnessed during the length of the test. The centre of mass of sphere A, which is attached to the source, appears to move in the first few minutes of the test (possibly as the x-ray source warms up), and then stays stable compared to the other spheres. This stability indicates that the source-detector system does not suffer from strong differential displacements that would cause an image to appear to displace.

The other two spheres, however, show relatively large displacements relative to the source-detector system. The tripod sphere (B) shows that the floor of the cabin is not at all stable relative to the source-detector system. Sphere C is attached to the rotation stage, which represents the position of the specimen relative to the source-detector system, also undergoes displacements throughout the test, although less intensely than the sphere attached to the floor. Such displace-

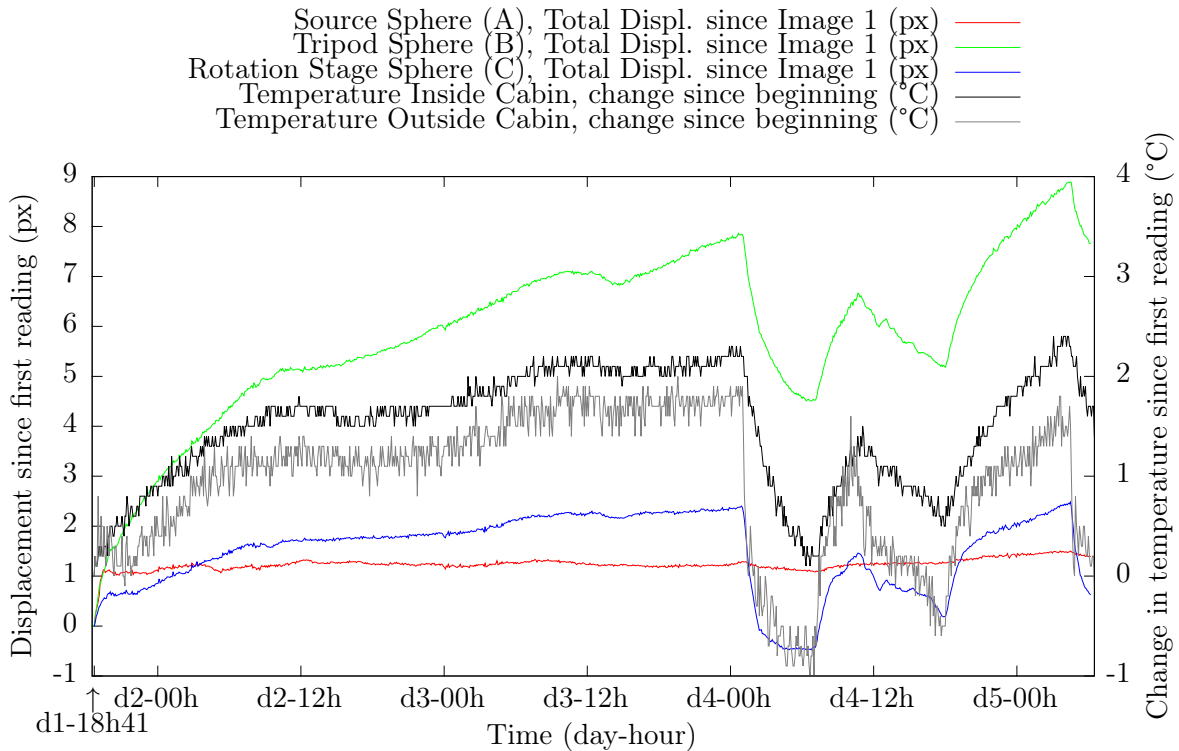


Figure 2.10: Results of temperature and displacement measurements made during the preliminary test

ments cause problems when reconstructing data, since the effective centre of rotation displaces during and between scans. The displacements of spheres B and C are very strongly (almost surprisingly) correlated with the ambient temperature inside the cabin.

The conclusions to draw from this preliminary test, despite its rather rough nature, is that hypothesis 1 appears to be pertinent only at the beginning of a scan, meaning that this problem can be avoided by small warm-up time before scanning. Hypothesis 2, that changes of temperature are responsible for differential displacements, seems to be relatively well verified from this test, insofar as the temperature is correlated to the displacements observed. The best candidate for root cause of the differential displacements between the rotation stage and imaging system due to temperature is the trolley on which the rotation stage is mounted (see Figure 2.3). Since the trolley is mounted in a cantilever-type arrangement on only one side of the cabin, it is quite possible that the strut that connects the horizontal surface (on which the rotation stage is mounted) to the vertical one that runs along the rails in the back on the cabin, flexes with changes in temperature. The effect of temperature was therefore studied further with a specially-developed piece of equipment: the TomoFrame[®], which is described in following section. This device was patented (which explains the late release of this PhD) and is also a registered trade mark.

2.3.2 TomoFrame[®]

The TomoFrame[®] was developed with Dr. Simon Salager in order to study differential displacements in the Laboratoire 3SR tomograph. The idea was to develop a tool capable of measuring the displacements of the trolley-and-rotation-stage system, relative to the imaging (source-and-detector) system, and also allow corrections of scans performed in conditions particularly vulnerable to differential displacements. To be able to correct radiographs prior to reconstruction, some known reference points need to be placed into each image, so it was elected to use the corners of the radiograph, since data acquired in the corners of the radiograph are partial (due

to the conicity of the beam) and therefore rarely used. The reference points were selected as 4 steel spheres, diameter 3.175mm. The rotation stage rotates during a tomographic scan, which means that a fixed reference point that appears close to the specimen in all the radiographs of the specimen cannot rotate with the specimen. Consequently the reference points must be mounted on the trolley directly - so that it moves with the rotation stage, but without rotating (assuming small thermally induced strain between the top of the trolley and the rotation stage). Due to geometrical constraints it was elected to design the TomoFrame[®] to be mounted after the specimen, in the direction of the beam (because for very zoomed scans, there is no room between the source and the specimen).

To best characterise the differential displacements thought to originate from thermal effects, an important requirement is to have a tool that is itself not vulnerable to thermal effects. Because of this, the TomoFrame[®] was designed to minimise the effect of internal thermal strains, so that its four reference points remain as fixed as possible in space relative to each other, for any likely change in temperature.




Figure 2.11: a): Schematic of the TomoFrame[®] (done by Gilbert Vian, Laboratoire 3SR) depth (not shown) is 8mm. b): Photograph of finished product with zoomed inset. The combination and order of the materials used is a key part of the design. c): Schematic of the TomoFrame[®] as intended to be installed in the tomograph (done by Gilbert Vian, Laboratoire 3SR)

The principal design principle behind the TomoFrame[®] is that the thermal strain induced by changes in temperature compensate each other. Simplifying the frame into a uniform 2D system, and assuming that thermal strains are a linear function of the change in temperature, consider the frame, as shown in 2.11 a), without the foot below. A positive change in temperature applied to this frame will cause the four corners of the TomoFrame[®] to move away from each other. The centre of the frame, however, will remain the centre of the circle which is shown on the figure. If the “legs” on which the spheres are attached are now considered, as the corners move away from each other, and from the centre of the shape, the legs will strain in the opposite direction,

i.e., towards the centre of the frame. They are not long enough, however, to compensate for the thermal strain of the entire frame, meaning that spheres in total still move outwards for a positive temperature change.

The thermal strain in the legs can be made to equilibrate the thermal strain in the frame by using a material in the legs which has a higher coefficient of thermal expansion than the frame. Therefore, to minimise thermal dilation effects in general, the main frame was built in Invar, a type of steel with a very low coefficient of thermal expansion ($1.5 \times 10^{-6}/\text{K}$ in this case, ordinary steel for reference is around an order of magnitude more). Most of the legs, as can be seen in Figure 2.11 are in fact contiguous with the frame, and so are also made of Invar. The thermal strains are equilibrated by an 0.4 mm hollow shim of aluminium (machined from a tube), with a coefficient of thermal expansion approximately equal to $23.1 \times 10^{-6}/\text{K}$, placed between the ends of the legs and the sphere. The use of aluminium here has the added advantage of being relatively transparent to x-rays, allowing the spheres to be easily identifiable in the radiographs.

In principle the design of the TomoFrame[®] could be scaled up or down in order to be the right size for any scan – an executive decision was made with this version, that the scans with the highest zoom would be targeted (which is why the TomoFrame[®] has not been applied in the scans performed in this work – it appears too small) since these have both the highest amount of zoom (allowing precise measurements to be made at least in the plane of the detector) as well as the largest amount of problems due to geometrical instability.

Figure 2.12 shows a radiograph from a 5 hour test scan on a specimen of Hostun sand in a state of partial saturation at $7.9 \mu\text{m}/\text{pixel}$ size. The TomoFrame[®] can be seen to occupy a little too much room in the radiographs in these settings, but has been adjusted to avoid blocking the view of the specimen. The four spheres can be easily identified and the aluminium shims nicely do their secondary job of separating the spheres from the frame in the radiograph, which facilitates automatic treatment. Unfortunately during this test no temperatures were recorded.




Figure 2.12: Projection from a scan on an unsaturated medium with the TomoFrame[®] in place, and visible on the outside of the specimen. During this scan the specimen rotates and the TomoFrame[®] stays relatively stationary (moving slightly in these images due to the differential displacements whose job it is to measure)

To correct the projections acquired with the measurement of relative displacement given by

the TomoFrame[®], the relative displacement of each radiograph from the first radiograph needs to be measured. This is achieved by isolating the four spheres in each image, and measuring their position (from their centre of mass) and their area.

The variation in the positions of the centres of mass of the spheres allows the measurement of the relative displacement of the TomoFrame[®], and consequently the specimen, in the plane of the detector, as well as the rotation in this plane (around an axis of rotation – see Section 5.3.4 – which is the axis of the beam).

As a tool for the study of thermally-induced strains, it is desirable to measure the entire 3D rigid body transformation (displacements and rotations) of the TomoFrame[®]. It is for this reason that the areas of the spheres are measured in each image, they can be used to calculate the radius of the spheres in the image, which can be used to calculate each sphere's position in the z (zoom) axis. There is a hypothesis made in this case that the stand which connects the frame to the trolley has very little thermal strain – this is because no temperature reading has been made of the cabin, or more specifically of the stand, so the small variations in vertical displacement cannot be accounted for. Since this measurement is a little noisy, it is fitted in order to smooth it. The displacement of the 4 spheres from their initial position in each radiograph is known at this point. These displacements are put into a least squares solver which tries to find the best mapping between the initial position of the spheres to each subsequent position with a 3-component translation vector and a 3D rotation matrix.

Figure 2.13 shows the evolution of these measurements during the 5 hour test-scan. The rotations measured are very small indeed, surprisingly, the angle with the smallest variation is the angle describing a rotation in the plane of the detector. This is quite possibly due to the fact that the z-direction of the spheres (coming from a measurement of the sphere's area) is measured much less precisely than their x and y-position (in the plane of the detector) – it is worth noting that in Figure 2.13, the z-position is given in mm, and not pixels. The largest rotation angle is around 0.02 degrees, which when applied in the middle of a radiograph causes the pixels on the far border to move around 0.33 pixels, which remains very small. The displacements, however – especially in the horizontal direction (on the detector) – show large variations. It is important to note that for the horizontal and vertical displacements must be scaled by the difference in the zoom between the TomoFrame[®] and the specimen (since the TomoFrame[®] is behind the specimen, and so benefits from less zoom, otherwise said, the same 1mm displacement will make the specimen move a lot in the image, and the tomoframe less, because of the difference in zoom or z-position). In this case the scale factor is 5.58, *i.e.*, a 2 pixel displacement of the TomoFrame[®] in the direction of the plane of the detector corresponds to a 2×5.58 pixel displacement of the specimen – in the acquired radiographies. Before the TomoFrame[®] was put in place, the only method of knowing whether differential displacements had taken place was to compare the first and last images acquired in the test because they are at almost the same angle, but at very different times. If an offset were measured (as it would be in this case, particularly in the x-direction) this could be applied proportionally to each radiograph in order to correct them, by making a hypothesis on the linearity in time of the observed displacement. In this case, approximating the vertical and horizontal displacements by a straight line going from their first to their last point would produce fewer errors than hypothesising that no movement has occurred at all, so a correction of this kind, for this particular test, would probably be relatively successful. It is evident, however, that this very strong hypothesis cannot be relied on in all cases.

A reconstruction using the raw radiographs is shown in Figure 2.14 a). This reconstruction has been made by finding the best centre of rotation with the tool described in Figure 2.8. It is immediately visible that this reconstruction is not satisfactory. The relative displacements and rotations of the TomoFrame[®] (in the plane of the detector) can be used to correct each radiograph's incremental displacement since the first radiograph. This is done by applying the measured rotation (in the plane of the detector) directly to the image, however the translation in the plane of the detector needs to be scaled (as stated above) from the TomoFrame[®] to the

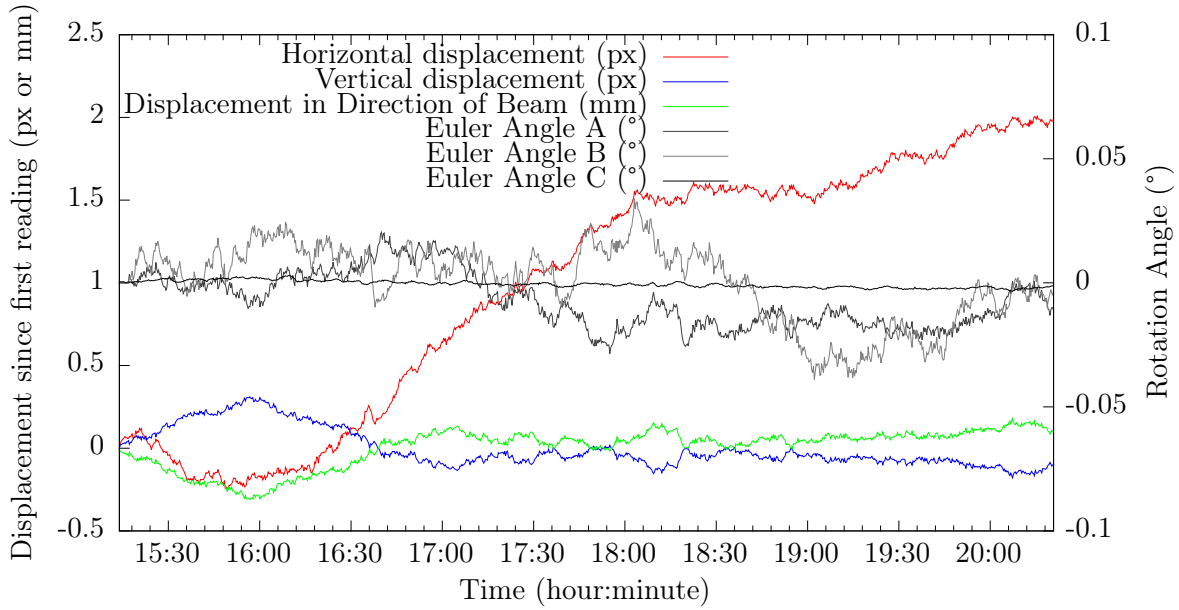


Figure 2.13: Graph showing the evolution of the rigid-body displacement of the TomoFrame[®] as calculated taking into account all four reference points, during a long scan

specimen, in order to take into account the difference in zoom. Radiographs are translated and rotated by bi-linear interpolation (modified from the 3D interpolator developed in Section 5.3.6), and are then for reconstruction. The results of a reconstruction (again using the best centre of rotation as described above) using these corrected radiographs is shown in Figure 2.14 b), and can be seen to greatly improve image quality.

Radiographs could potentially be corrected for displacements in the direction of the beam (*i.e.*, changes in zoom), but although these are measured, this has not yet been implemented. The other two components of rotation, however, cannot be directly applied to correct an image (one of the angles, for example describes a rotation around a vertical axis, such as the one that the specimen is rotating around during the tomography. It is clear that this rotation of a specimen cannot be corrected by a rigid-body displacement on the recorded radiograph). All the measurements of the rigid-body transformation of the specimen, as measured by the TomoFrame[®], could however potentially be taken into account at the reconstruction phase (in particular in algebraic reconstruction methods).

Differential displacements: conclusions and further work

A tool for the measurement of the rigid-body transformations of the rotation stage has been developed, and appears to work well both for the measurement of transformations, as well as in the correction of radiographs acquired with the tool in place (see Figure 2.14). A further interesting initial result from the first test with the TomoFrame[®] is that the rotations measured appear to be very small. If after further studies rotations are systematically small, it may be sufficient to include a single sphere in subsequent high-zoom scans. This would allow the tools developed for correcting projections to be used to apply displacement corrections (coming from measurements on a single sphere) to acquired radiographs.

Work is ongoing towards a fully instrumented test, with several thermocouples and with the TomoFrame[®] in place, with the objective of better characterising the relationship between fluctuations in temperature and fluctuations in differential displacement between the rotation stage and the imaging system. This may, in the future, lead to changes in the tomograph's design in order to stabilise it against such temperature dependent differential displacements.




Figure 2.14: Slices from reconstructed volume of a partially saturated sand with and without correction for differential displacements

2.4 Conclusions

This chapter has shown some important details about the physics of x-rays as they apply to imaging of geomaterials. The setup in Laboratoire 3SR used to acquire the images presented herein has been described, as well as the conditions and settings in which the scans have been made in this work. Subsequently, the procedure for reconstructing a 3D field of x-ray attenuation in the scanned domain was detailed. The 3D images coming out of the reconstruction process are the input for chapters 4 and 5, and will be used to make micro-mechanical measurements on the deforming media.

The stability of the x-ray acquisition system is currently being studied with some specially-designed tools, preliminary findings show a strong correlation between temperature and differential displacements between the rotation stage and the acquisition system. The “TomoFrame[®]” has been developed in order to put four reference points (which have been designed not to move relative to each other even with a difference in temperature), into the corners of radiographs acquired during a tomography.

Chapter 3

Experimental campaign

This chapter details the mechanical tests performed in this doctoral work. The previous chapter has detailed how 3D images representing fields of x-ray attenuation are obtained from x-ray tomography. Since x-ray tomography is a non-destructive imaging technique, several images of a specimen can be acquired with minimal disturbance of the specimen. The objective of this work is to analyse the *evolution* of the micro-structure of (some) granular media during deformation. The following chapters will detail the data treatment required to pass from images to measurements of micro-structure. This chapter will describe the experimental procedures and mechanical tests performed. A principal objective of this work is to make observations at the grain scale, specimen sizes have therefore been greatly reduced from standard soil mechanics sample sizes to be able to take advantage of the zoom of the x-ray scanner's cone beam so that individual grains are visible in the tomographic reconstructions. Triaxial compression has been selected as the mechanical test which will be performed on specimens of the granular materials being studied. The specimens used in this work consequently are cylinders of approximately 22 mm height and 11 mm diameter. Triaxial compression tests are widely used standard, however the tests described in this work are highly specialised, differing principally (besides in specimen size) in the measurement instrumentation: the tests are performed *in-situ* inside the x-ray scanner in Laboratoire 3SR. The tests described in this work are far from being part of standard experimental geomechanics practice.

A range of different materials has been tested in this work. However, the analysis will focus on a series of tests performed on three different materials: Hostun HN31 sand, Ottawa 50/70 sand, and Caicos ooids. Each material has been deformed in triaxial compression at 100 and 300 kPa confining pressure (with one test at 400 kPa for Caicos) *in-situ* in the x-ray scanner.

In this chapter all tests are drained $\sigma' = \sigma$, so all stress quantities will simply be denoted by σ .

Layout of chapter

This chapter introduces the triaxial test in Section 3.1 and then details the materials used in this work in Section 3.2. The chapter then goes on to detail the tests performed in this work, by first covering the method used to prepare specimens (Section 3.3), followed by a description of the equipment used (Section 3.4) and finally the way in which the specimens are tested in Section 3.5.

The experimental campaign is detailed in Section 3.6 by describing the mechanical responses of the different materials tested. This section shows the axial stress and volumetric strain *vs.* axial strain responses of the specimens tested. Cross-sections from 3D images of each specimen in its initial condition are also shown, evaluate each specimen's initial geometry. These macroscopic responses are commented, for each test.

The representativity of the small specimens size and associated testing are discussed in detail in Section 3.7.

3.1 The triaxial compression test

The triaxial compression test is a standard test in soil mechanics and has been chosen in this work for two principal reasons. First, the use of a standard test allows easier comparison to existing results. Second, the cylindrical shape of the specimens is very convenient for x-ray tomography (since objects are rotated around an axis and the shape of the domain with full information is very close to a cylinder for a cone-beam setup – and exactly a cylinder for a parallel beam).

A complete reference for the technical details of the triaxial test is Bishop and Henkel (1962). Several international standards also exist for performing a triaxial test, however these only have paid-for access:

- ASTM WK3821 (2011). New Test Method for Consolidated Drained Triaxial Compression Test for Soils (under development). ASTM International, West Conshohocken, PA
- ISO/TS 17892-9:2004 (2007). Geotechnical investigation and testing – Laboratory testing of soil – Part 9: Consolidated triaxial compression tests on water-saturated soils. International Organization for Standardization
- BS 1377-1 (1990). Methods of test for soils for civil engineering purposes. General requirements and specimen preparation. BSI. ISBN 0-580-17692-4

Very simply, a triaxial test on dry sand (like the ones performed in this work) is performed on a cylindrical specimen (typically, but not necessarily, of slenderness ratio 2 which offers a compromise between end effects and propensity to buckle) prepared in a flexible membrane that is held in a stretcher or mould during specimen preparation. The cylinder’s axis is vertical at the beginning of a test. On the top and bottom of the cylindrical sand specimen, metallic loading platens are added. Several different standards exist for the dimensions of the cylinder; in Laboratoire 3SR specimens of 100 mm diameter are generally used. The specimen is tested in a triaxial cell, which is filled with a confining fluid and then pressurised. The specimen is then deformed in the direction of the axis of the cylinder by advancing a ram against one of the loading platen. The ram is typically advanced under displacement control, which allows exploration of the post-peak response of the specimen. The shortening imposed on the specimen is measured along with the specimen’s reaction force. The contact between ram and specimen can be defined either in such a way that the loading platen that is displaced is always kept horizontal and axially centred, or it can be allowed to rotate by applying load with a point contact. Both options have their advantages, however for this work the latter has been selected since it does not carry moments, and therefore allows more freedom to the specimen to deform.

Between the end of the specimen’s preparation and the application of the cell pressure, suction is typically applied to the inside of the specimen so that there is a pressure difference over the specimen membrane, which applies a mean effective stress to the specimen. This is done so that the specimen’s granular arrangement does not change when the specimen is removed from the mould: the effective stress allows the specimen to stand under its own weight (plus the weight of the top platen).

During a triaxial compression test, the granular skeleton initially undergoes isotropic compression due to the vacuum applied to the inside of the specimen. The isotropic compression is continued while the cell pressure is increased (at the same time the pressure inside the specimen is slowly increased to atmospheric) up to the test pressure. Thereafter the specimen is sheared by applying deviatoric loading, σ_1 is increased (by the ram), while $\sigma_2 = \sigma_3$ (the cell pressure). This means that the triaxial test’s name is slightly misleading since not all three principal stresses are controlled independently (machines that allow said control are called “True Triaxial”).

When analysed in the q/p (deviatoric/mean stress) plane, the isotropic portion of the test is simply a horizontal line up to the p value of the cell pressure. The deviatoric loading stage represents a line of gradient 3 since in these conditions $q = \sigma_1 - \sigma_3$ and $p = \frac{\sigma_1 + 2\sigma_3}{3}$.

3.1.1 Specimen size

The objective of this work is the experimental characterisation of the micromechanics of deformation in granular materials. In light of this, it is highly desirable to perform x-ray scans at a spatial resolution such that individual grains can be identified, in the resulting 3D image. The combination of the fixed resolution and physical size of the detector, the zooming effect of the x-ray cone beam and the particle sizes to be studied impose an upper limit on the size of the specimen that can be tested while being able to image all the grains of the specimen in 3D – this is detailed in Section 2.2.1.

A pixel size around 15 μm has been selected in order to best image the granular materials being studied, this limits the field of view to around 30.5×23.0 mm. Due to the conicity of the beam, the top and bottom edges of the radiograph only have partial information, a specimen size equivalent to the specimens created for previous work in the ESRF is selected (see Hall *et al.*, 2010). The specimens used in these *in-situ* tests, are therefore much smaller than standard: they are a little under 11 mm in diameter and aim to be between 22-24 mm in height.

The membrane allows the confining pressure of the triaxial cell to exert a pressure on the specimen, since the pore pressure inside the specimen is always the atmospheric pressure. Membranes for each test are cut from made-to-order tubes, which are 10 mm wide and 0.3 mm thick.

3.2 Materials tested

In this work three granular materials are studied: Hostun sand, Ottawa sand and Caicos ooids. These materials cover a range of different grain angularities and are all poorly-sorted. This section describes the three materials.

3.2.1 Hostun sand

Hostun sand is produced in a quarry close to the commune of Hostun in the department of Drôme, in the Rhône-Alpes region of France. It is approximately 60 km from Grenoble. The sand tested in this work was collected directly from the quarry, which is run by Sibelco France, at the beginning of 2011. The material is excavated as clumps made of a range of different silica grain sizes, lightly cemented by kaolin. The silica particles making up these clumps have been crushed *in-situ* from larger rocks, and have consequently not been transported very far. Since particles have been generated by crushing and have not been transported very far they are very angular. This can clearly be seen in Figure 3.1, which shows an SEM image of Hostun sand grains (see Flavigny *et al.*, 1990). Hostun sand grains are also visibly angular in the x-ray tomography images acquired in this work, as can be seen for example in Figure 4.1. As a last example, Figure 4.13 shows a radiograph from a nano-tomograph of a single Hostun sand grain showing again large angularity.

After quarrying, clumps are washed in order to remove the kaolin in a machine called a “Hydrocyclone”, and are then sorted by grain size in a “Hydroclasseur” (“watersorter”), which is an expanding flume, where the heaviest grains collect at the beginning of the flume, where the flow is the fastest, whereas smaller particles travel further. This explains the tight grain size distribution shown in Figure 3.4, calculated for Hostun sand from data provided by the manufacturer Sibelco France (2011). Therein the D_{50} is given as 338 μm .

The sand is used in this work is known and marketed as “HN31”. This sand is the successor, and likely to be practically identical to the “S28”, and “RF” varieties that have thoroughly studied




Figure 3.1: Scanning Electron Microscope image of a few grains of Hostun sand from the PhD thesis of Colliat-Dangus (1986) and published in Flavigny *et al.* (1990)

in Grenoble. The mechanical behaviour of S28 and RF sands were compared in Combe (1998) and found to be equivalent. RF and S28 varieties of Hostun sand are quarried either side of the current quarry for HN31, both within a kilometre.

3.2.2 Ottawa sand

Ottawa sand is the sand of intermediate angularity studied in this work. It comes from sedimentary deposits in Ottawa, Illinois (U.S.A.) and was kindly provided by Jason DeJong from UC Davis. The variety of Ottawa sand studied is known as Ottawa 50-70 sand (sometimes Ottawa 50/70 sand), which means that it has been sieved between the 50 and 70 US Standard Mesh sieve sizes, corresponding to mesh sizes of 300 μm and 210 μm respectively. Ottawa 50-70 sand is a relatively well-studied sand, by various groups such as Georgia Tech and the University of California, Davis; however there seems to be a considerable variability in the properties reported in the publications about this sand. For example, the D_{50} values reported in literature vary considerably, but generally (though not always) fall with the 210-300 μm range. It has been estimated that the D_{50} of the sand used in this work is in the order of 250 μm . In particular Mortensen Montoya (2012) – a PhD thesis recently defended in UC Davis, where the sand tested in this work come from – is used as a reference for this sand. The D_{50} reported therein is 220 μm – however no grain size distribution is reported. The data from the grain size distribution is therefore taken from Kim and Santamarina (2008).

Like Hostun sand, Ottawa sand is a siliceous sand made up of quartz grains, but due to the deposition process that it has undergone, grains are more rounded, as can be seen in Figure 3.2 – particularly when comparing it to Figure 3.1.

3.2.3 Caicos ooid

Caicos ooids are a much rarer material than Hostun sand and Ottawa sand, which are well known and easily available. The one litre bottle of Caicos ooids kindly provided to Laboratoire 3SR by Hubert King of Exxon Mobil was given as “the last bottle” of this material, and so where possible, the material has been used with reserve. More material can, in principle, be collected however no guarantee can be made of the similarity of newly collected material. The material




Figure 3.2: Scanning Electron Microscope image of a few Ottawa sand grains (J. DeJong, personal communication)

comes from Ambergris Shoal in the Caicos platform (Turks and Caicos Islands), in the British West Indies.

Unlike the two sands described above, which come from the dissociation of larger rock masses, ooids grow from small seeds such as shell fragments or small quartz or calcite grains, and progressively become larger. Ooids grow in marine environments, where the waves drive the accretion of material around the seed. Two principal mechanisms exist for an ooid to grow: by physical attachment of material as the ooid rolls (like a snowball), or by precipitation over the surface of the ooid. Both these mechanisms tend to generate rounded grains – from whence the name comes – although in the images acquired, some grains which are prolate spheroids (“rugby ball”, close to cylinders) have been seen.




Figure 3.3: Radiograph from a nano-tomograph of a single Caicos ooid, diameter approximately 500 μm . Courtesy of RX-Solutions, Annecy

Figure 3.3 shows a *radiography* of a single Caicos ooid of the ones tested in this work. An ExxonMobil internal study of this material using x-ray diffraction, which has been shared with Laboratoire 3SR, reveals that these grains are made of CaCO_3 and that more than 96% of the

material is Aragonite, the rest being calcite and “high magnesium calcite”. At the edges of the radiography shown in Figure 3.3, zones of porosity can be seen, which corresponds to an observation from ExxonMobil that an outside shell of these grains, 3 to 10 μm thick has porosity around 15%. Some grains have been observed to have internal porosity in x-ray tomography images both by ExxonMobil as well as in the context of this work.

The ooids studied in this work have been cleaned and sieved, and are thus also poorly graded, as can be seen in Figure 3.4 – this grain size distribution has been prepared with data communicated by ExxonMobil. Some grain-sized shells have survived the cleaning, and thus are occasionally visible in the specimens imaged – Figure 4.1 shows a slice from a 3D image where a shell is visible. The D_{50} of this material is around 420 μm , which makes this the biggest of the three materials studied.

3.2.4 Conclusions

The three materials studied cover a wide range of the angularities found in natural granular materials. The Hostun sand (most angular) and Ottawa sand (intermediate angularity) are both silica sands generated by weathering and breakage, whereas the Caicos ooid is a carbonate sand that has grown by rolling and is consequently very rounded. The different genesis of the ooid gives it some unusual properties, such as internal porosity in some grains, as well as a more porous outer shell.




Figure 3.4: Grain size distribution curves for the three materials studied in this work. Data for Hostun sand comes from the manufacturer (Sibleco France, 2011), data for Ottawa sand from Kim and Santamarina (2008), and data for Caicos ooids were supplied by ExxonMobil in a private communication

Figure 3.4 shows the grain size distribution for the three materials studied, all of which are poorly graded. Table 3.1 outlines the principal characteristics of these materials.

Name	Hostun HN31 Sand	Ottawa 50-70 sand	Caicos Ooids
Grain mineralogy	Quartz	Quartz	CaCO ₃ (mostly Aragonite)
D_{50}	338 μm	310 μm	420 μm

Table 3.1: Table showing some of the principal characteristics of the materials studied in this work

Figure 3.5 shows histograms of grain volumes (in voxels, which represent a cube of volume approximately $15 \times 15 \times 15 \mu\text{m}$ in volume). The histograms presented show that the grain volume distribution as measured by x-ray tomography for Hostun and Caicos is similar, with

Ottawa sand being considerably smaller. This has a slight discrepancy with the grain size distributions presented in Figure 3.4 insofar as the D_{50} of Caicos ooids appears larger than Hostun sand. This may be due to the fact that the sieve analysis performed in ExxonMobil may have been on a slightly different material, or that the material has in some way evolved in time, or with transportation.

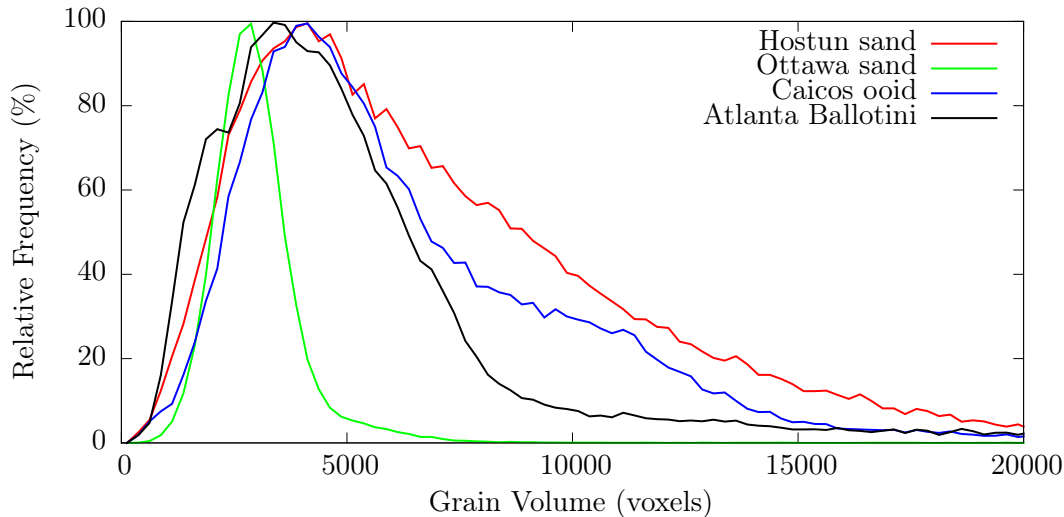


Figure 3.5: Grain volume histograms from tests HNEA01, OUEA06, COEA01 and ABEA02

3.3 Specimen preparation

The small size of the specimens means that the preparation of these specimens differs from the ordinary standard adopted in experimental soil mechanics. The preparation of such small specimen is of considerable technical difficulty, and consequently the quality and reproducibility of these specimens is a major concern. The quality of a specimen includes parameters such as geometric quality (how close the final specimen is to a cylinder of slenderness 2) and material quality (the uniformity of the distribution of the specimen's properties).

Specimens are prepared by air pluviation and tested dry (in drained conditions). The contact between porous stone and granular medium is not lubricated.

The choice of dry tests simplifies the means that a specimen does not need to be saturated in water. If the pore space of a specimen is not saturated, the volume change of the specimen cannot be measured directly. However, this is not a big hindrance, since the x-ray scans performed at several points during the test allow the reconstruction of an entire 3D image of the specimen, from which its volume can be calculated (this is shown in Section 4.4). A dry test also means that the bulk density of the specimen is lower than if the specimen were saturated with water. A reduced bulk density due to a less dense interstitial pore fluid is beneficial for two reasons: first, x-rays can penetrate it more easily, thus reducing scanning times and second, the contrast between solid and pore phases is increased.

3.3.1 Preparation of membrane for pluviation

The preparation of a specimen of a cohesionless granular material requires an initial arrangement of grains to be created inside a membrane. This differs from the preparation of specimens of materials with cohesion, such as clay, which already have an initial arrangement which holds steady even without confinement. This means a membrane can be stretched around these specimens with relative ease (once they are prepared). Since (by definition) a specimen of a cohesionless material cannot support its own weight without effective stress, a given initial arrangement of

grains cannot simply be created and then encircled by a membrane as with clay. As stated in Bishop and Henkel (1962): “In order to make a sand specimen for use in the triaxial test it is necessary to use a former [mould] which will maintain the required specimen shape until effective stresses of sufficient magnitude to make the specimen self-supporting can be applied.”

Figure 3.6 shows the way in which a specimen is prepared in this work. Figure 3.6 a) shows the holder on which the specimen will be created. The part of the holder which is visible is a hollow steel tube, which will allow suction to be applied to the inside of the specimen, when it has been created. A porous stone (diameter 12 mm) is then added to the holder in Figure 3.6 b). Two different types of porous stone have been used in this work, and lubrication was not added in both cases. The second type of porous stone was selected for its lower x-ray attenuation coefficient, which makes it easier to remove it from the acquired images – see for example OUEA04 compared to the other tests on Ottawa sand in table 3.4.

The porous stones have a 500 μm hole drilled axially through them – to facilitate the application of suction to the inside of the specimen. Specimens are dry, and there is consequently no exchange of liquids, the porous stone is not being used for its porosity, but rather for its density, it allows the top of the specimen to be imaged without completely blocking the x-rays as a steel surface would. Once the porous stone is in place on top of the specimen holder (Figure 3.6 b), a membrane is cut to size and slipped over the system so that it completely covers the porous stone and covers the holder to at least 3 mm in order to stop water invasion when the system is submerged. As can be seen in Figure 3.6 c), the membrane is actually stretched over the holder and porous stone, since its radius is 1 mm smaller than the holder and porous stone; it is kept in place by friction with the specimen holder. Great care is taken to ensure that the membrane is both vertical (as shown in the image), as well as well-centred with respect to the porous stone. The position of the membrane is regulated by hand by pulling the membrane up or down around the holder and inspecting the position of the membrane by eye. Note that the centring of the membrane can be rendered more difficult by slight inequalities of the thickness of the membrane.

In a typical triaxial setup, a rubber O-ring would be put over the membrane, pressing against the holder, to better guarantee the seal protecting the inside of the specimen from invasion of the confining fluid. O-rings are not used in this case because preliminary tests showed that the action of adding an O-ring causes the membrane to slip, and therefore for the membrane to lose its verticality.

Figure 3.6 d) shows an interim stage of the installation of the mould designed specifically in the context of this work – incorporating a previous design used for previous tests (of the same dimensions) at the ESRF in Grenoble. The mould used in this work uses the thicker part of the holder as a guide, and the change in radius of the holder as a support (in order not to press down on the membrane). The mould as shown has an internal cylindrical hole inside which the membrane passes. The mould, when closed (as shown in Figure 3.6 e) is airtight at the bottom, however access to the inside of the mould is possible by a horizontal nozzle (shown with a tube attached in the photo).

Figure 3.6 f) shows the membrane being folded back over the mould. Once this is done, the air which is between the mould and the outside of the membrane is removed by a vacuum pump connected to the tube shown in the photo. Care has been taken (before closing the mould) to place a piece of thin tracing paper or gauze over the point at which the hole enters the inside of the mould. This is done to avoid the membrane being sucked into the hole and stopping the vacuum forming all the way around the membrane. Application of a vacuum causes the air to be removed between the outside of the membrane and the mould, pushing the membrane against the mould, so that it adopts its shape, which is the desired shape of the outside of the specimen.

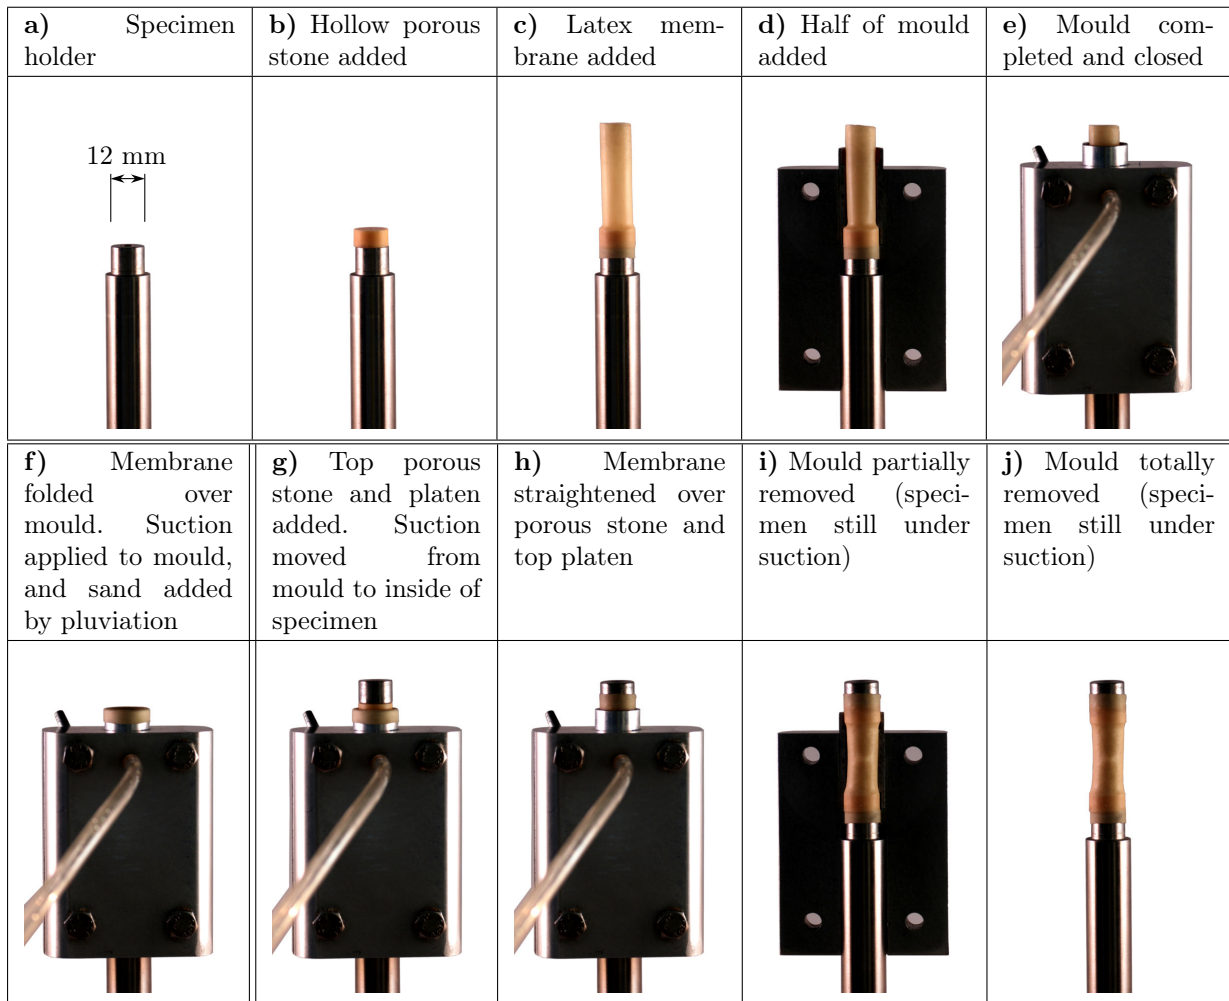


Figure 3.6: A series of photographs illustrating the procedure by which a specimen is prepared in this work. Between frames f) and g) the material is pluviated into the membrane

3.3.2 Deposition of grains into membrane

Once both sides of the mould are closed around the sample, the top of the membrane is folded over the mould, allowing access to the inside of the membrane from the top. The sides of the membrane are pushed against the mould by the applied vacuum. A specimen of granular material can now be created in this space inside the membrane.

Dry pluviation has been selected in this work as the method for the deposition of the grains into the membrane. Dry pluviation consists in the dropping of grains into the specimen, so that grains accelerate due to gravity. The fact that grains are travelling at speed means that when they come into contact with the other (stationary) grains that make up the specimen, they will have the energy to pack themselves in a dense fashion. The amount of energy, and thus the density of packing can be varied by changing the drop distance. This specimen preparation technique has been selected for its reproducibility, since other techniques for achieving specimens (be it dense or loose) generally require interaction with the specimen during its preparation (*e.g.*, moist tamping), which is extremely difficult with the dimensions of these specimens. These techniques are likely to be much less scalable to such a small specimen.

Once the membrane is stretched in the mould, the specimen is “pluviated” into the membrane. In the preparation of standard triaxial tests, the granular material is typically poured into a tube that ensures that grains fall principally in a vertical direction. For standard triaxial test the tube

is considerably smaller than the specimen, and so the specimen is created with the nozzle or end of the tube *inside* the membrane. This was however not deemed possible for the specimens in question because of the relative sizes of the grains to be deposited and the diameter of the specimen. A tube small enough would have had to have such a small diameter when compared to the width of the specimen, that grains would likely bounce from wall to wall during their descent, thus causing grains to be deposited in a biased and less-dense fashion. Therefore, a tube of length 108 cm (with a funnel width 125 mm installed on the top to direct the flow of grains) with a width of 25 mm is used to preserve the verticality of the trajectories of the grains. However, this means that many of the grains poured into the pluviation system do not make it into the specimen, and simply bounce off the mould. This is not a problem in terms of the volumes of granular materials used, with the notable exception of the Caicos Ooid material. In the case of Caicos Ooids therefore, this system of pluviation was used only once to make a specimen (COEA01), and thereafter pluviation was done at a much smaller height with a cone-with-a-hole, made from paper with a aperture of approximately 20° and a hole for pluviation 4 mm in diameter.

Both systems for pluviation are used in such a way that, as the specimen is created, the pluviation system is lifted at approximately the same rate as the advance of the top of the sample, to ensure a constant drop height.

The pluviation of grains is stopped slightly after the final height is achieved (a slightly too-high specimen is created), and then extraneous grains are removed with a sticky, flat surface to help ensure a flat and horizontal top surface.

3.3.3 Removal of mould

Once the desired granular arrangement is achieved (by pluviation of grains into the membrane) the specimen now needs to be completed and freed from the mould. A porous stone like the one at the bottom of the specimen, but for the fact that it does not have a hole drilled through it, is now added to the exposed top surface of the specimen. This is followed by a steel platen – on which the ram that applies the axial load will press. These two elements are placed on top of the specimen, as per Figure 3.6 g), and their horizontality is verified. If found to be unsatisfactory, the porous stone can be carefully removed, and some more grains removed from the surface of the specimen, so that the porous stone and platen can rest horizontally on top.

Once all the components of the specimen correctly in place, the vacuum is removed from the mould, which means that the membrane will no longer be held to the sides of the mould and will press on the granular skeleton. The membrane is then removed from around the mould and straightened over the assembly of porous stone and platen, as show in in Figure 3.6 h).

At this stage, the vacuum pump is connected to the bottom of the holder, thus applying a vacuum to inside of the specimen. The fact that the pressure in the voids of the specimen is below atmospheric means that the specimen has non-zero mean effective stress, thus it can now stand under its own weight (and the weight of the porous stone and platen that have been added to the top). At this point, the mould around the specimen is carefully removed – as can be seen in Figure 3.6 i) and j).

The specimen is inspected for its verticality. If not found to be acceptable, the procedure is started again from the beginning, without reusing the membrane.

It is important to note that the specimen will be tested upside down – that is to say that the holder on which the specimen has been prepared will be installed into the top of the cell, and the specimen will be suspended in the cell from the top, and then be compressed from below by driving *up* a ram against the platen on the free end of the specimen. See Figure 3.8 (“Sample holder”).

3.4 μ Tomo triaxial apparatus

This section details the triaxial system used in this work, in particular the cell and loading system used.

3.4.1 Loading system

In order to apply the deviatoric stress required for triaxial compression, the specimen is compressed axially by a ram driven up by the loading system which is described in this section. The objective of the loading system is to apply and measure the axial force compressing the specimen from below, as well as measuring the vertical displacement of the ram. These key measurements allow the macroscopic axial stress and axial shortening to be measured during a triaxial test.

Technical details

The loading system shown in Figure 3.7 consists principally of an electric motor with two gearboxes with a gear ratio of 1000 (developed and detailed in Lenoir, 2006). The rotation generated by the motor and the gearboxes is turned into a vertical movement by a worm drive which drives a loading head (shown in Figure 3.7) up or down. The use of such a large gear ratio combined with the use of a worm drive means that a fine control over loading velocity is achieved. The speed range for the loading head is from 0.1 $\mu\text{m}/\text{min}$ to 100 $\mu\text{m}/\text{min}$, which corresponds to a strain rate from 0.0005% to 0.45% per minute (for a specimen of height 22 mm). The maximum axial force that this motor and gearbox system can apply is reported as 7.5 kN. The motor is driven remotely from the laptop used for data acquisition.

Installation of loading system components

The entire loading system is suspended below the rotation stage by four tie bars (one of which is indicated in Figure 3.7). Between the cell and the loading system there is a circular plate (which cannot be seen in Figure 3.7 but is however shown in Figure 3.8) through which the tie bars also pass. The circular plate is responsible for making a seal both with the inside of the cell (with an O-ring) as well making a seal with the loading ram. Both of these seals are made in order for the cell fluid and consequently the cell pressure not to be lost. The two seals that the circular plate makes are visible in Figure 3.8. The ram is pushed progressively through this seal by the loading system. The seal between the ram and the circular plate was originally designed with floating O-rings, which were found to be too unreliable under pressure, so the O-rings were removed. In this work, rams are systematically greased to minimise leakages. The fine layer of grease is responsible for a small amount of friction as the ram moves through the seal – this was characterised (by driving the ram up and down in cycles) and subtracted from force readings.

The ram itself is a steel cylinder 10 mm in diameter, which has one flat end and one hemispheric end. The hemispheric end of the ram is the one in contact with the bottom of the specimen. A domed ram like this allows rotation of the bottom platen, and, since there is a very small area of contact, no moments are transferred.

Measurements

To measure the force applied by the ram onto the specimen, an HBM C2-500N force meter is installed on top of the loading head, in contact with the bottom of the axial loading ram. This piece of equipment was calibrated at the beginning of the work against a reference force meter in Laboratoire 3SR.

The measurement of axial displacement is made by an LVDT (shown in Figure 3.7), which is attached to a tie bar and measures the vertical displacement of the loading head.

Safety

The loading system has a set of proximity sensors that operate cut-outs on the motor to stop it driving the loading head too high (to avoid the force meter being squeezed between loading head and the circular plate).




Figure 3.7: Labelled photo (with background faded) of triaxial cell in place in the x-ray tomograph in Laboratoire 3SR

3.4.2 Triaxial cell

Compared to the designs presented in Bishop and Henkel (1962) reproduced in many laboratories, the triaxial cell used in this work has some supplementary design constraints to satisfy, due to the use of x-rays.

Since the overarching objective of this work is to perform *in-situ* x-ray scanning, the cell must be able to be scanned without blocking the x-rays. This means that the vertical steel return-bars or tie bars typically used to equilibrate the axial load imposed on the specimen cannot be used, since these would attenuate x-rays too much.

The job of the tie bars is performed (in this design) by the cell – the vertical force imposed on the specimen by axial loading are equilibrated by stresses in the cell itself, which is also sustaining the cell pressure. As described in Hall *et al.* (2010) and Lenoir *et al.* (2007), the cell is designed in PMMA Poly(methyl methacrylate), known commercially as Plexiglas or Perspex. It is a transparent plastic, which allows visual monitoring of the test, but, more importantly, is

also very transparent to x-rays. A trade-off has had to be made between the wall thickness of the cell, which allows greater loads to be carried, and the transparency to x-rays, which is improved by a thinner cell. This “low-pressure” cell has relatively thin walls (8 mm), and thus allows easy x-ray scanning, but is limited for example to 1000 kPa cell pressure.

The cell, visible in Figure 3.8, is I-shaped in order to be able to translate it very close to the x-ray source for very zoomed-in scans, although this has not been done in this work.

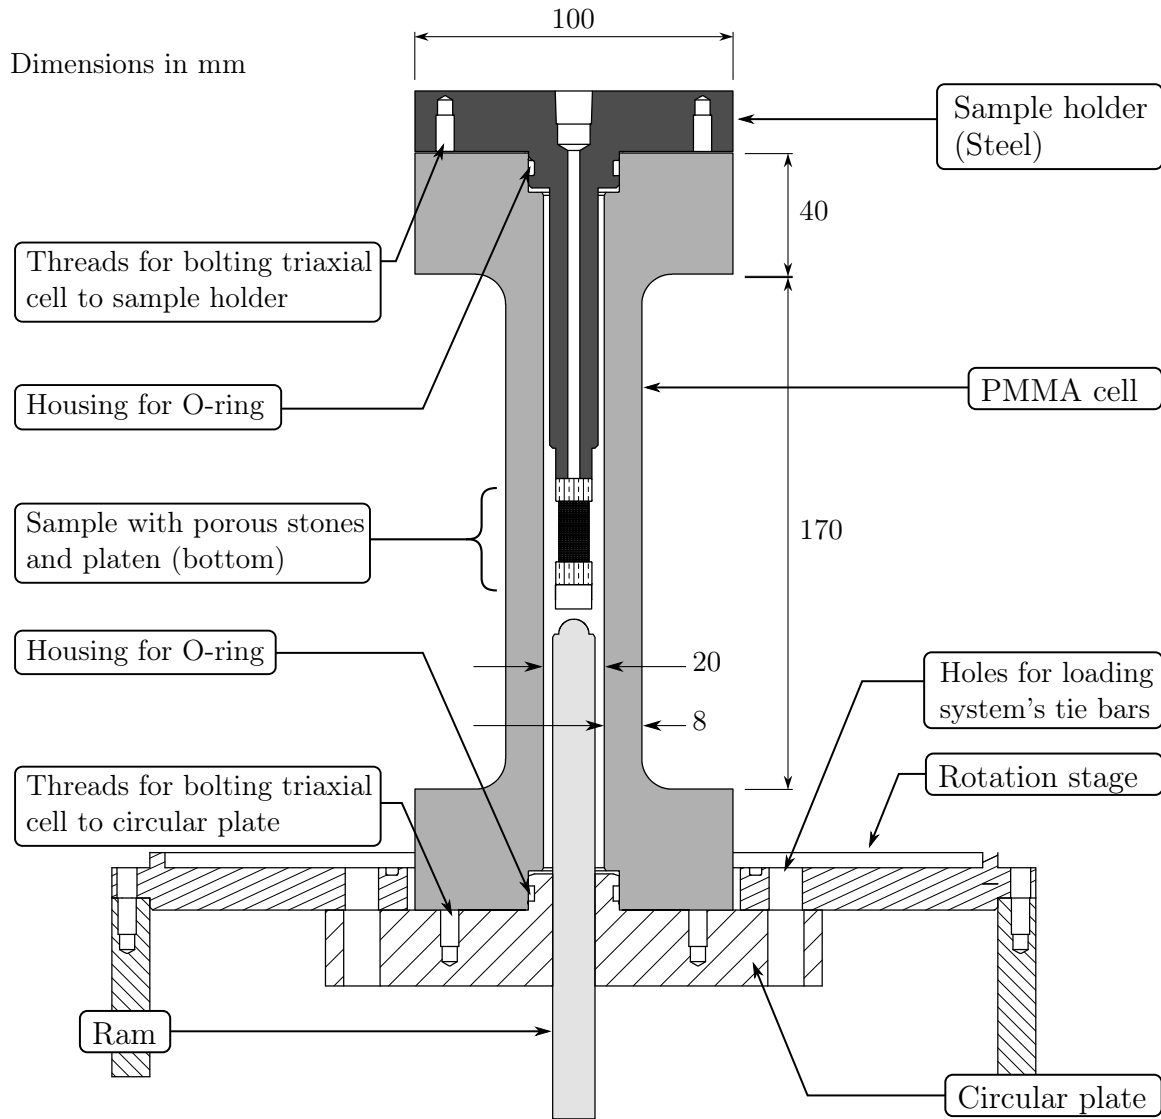


Figure 3.8: Schematic of triaxial cell created by Pierre Bésuelle and Nicolas Lenoir. Dimensions in mm

The cell has a circular hole of diameter 20 mm all the way through which expands to a wider 30 mm at the extremities. The bottom part of the cell has a tap on the outside that communicates with the internal space in the cell via a small, horizontal hole. The tap allows filling of the cell with confining fluid. As described above, in Section 3.4.1, the bottom part of the triaxial cell makes a seal with a circular plate, which has an O-ring that seals the bottom of the cell (where its central opening widens to 30 mm). The cell is also bolted into this circular plate (black bolts can be seen in Figure 3.7) in order to equilibrate the vertical forces carried by the walls of the cell generated by the pressurisation of the cell, as well as axial compression.

The top of the cell has a small, horizontal hole which allows communication with the inside of the cell from the outside. Compressed air is attached to this hole, to pressurise the pore fluid. Once a specimen has been prepared on the specimen holder, it is installed into the cell with the

top of the specimen (as prepared) pointing down. The specimen holder makes a seal with the top part of the cell.

3.5 Testing procedure

After a specimen has been prepared, it is installed into the cell – the test is then set up by preparing all the measurement equipment, and filling the cell with confining fluid (described in Section 3.5.1). Thereafter data acquisition is started and the triaxial test begins, as described in Section 3.5.2.

3.5.1 Test setup

Installation of principal components

When setting up for a test, the loading system is set up outside the tomograph’s cabin. The force meter is installed onto the loading head, and the LVDT is loosened in order to allow displacements of the loading head. Thereafter, the circular plate, which seals both the bottom of the triaxial cell and the ram, is slid onto the tie bars of the loading system. The height of the exposed ram is checked, and if too high (*i.e.*, not enough space is left for the specimen holder and the specimen), the loading head is lowered causing the ram to descend. This system is then installed under the rotation stage; the loading system’s tie bars pass through four holes in the rotation stage’s hollow plate and are secured by four bolts on the top side of the rotation stage.

Once the system is in place, the ram is greased, and slid into position downwards through the circular plate. The PMMA cell is then pushed down over the O-ring on the circular plate, then bolted into place.

Installation of the specimen

Section 3.3 has described the specimen preparation procedure. At the end of this procedure (see Figure 3.6 j), a specimen is created on the specimen holder, and suction is applied through the suction holder to the inside of the specimen, so that some effective stress hold the granular skeleton together. At this point (as briefly described above) the specimen holder is slowly turned upside down – with the vacuum pump still connected to the specimen holder – and is lowered vertically into the exposed top part of the cell. The holder is carefully tightened over the O-ring and the specimen is checked again for its geometry. If found to be slightly tilted, the holder is removed, and small adjustments of a few degrees can be imposed by deforming the specimen by hand. This is likely to cause local rearrangements of grains and to create some voids in these dense specimens, but these interventions are limited to small modifications of specimen geometry.

Preparation for mechanical loading

With the specimen in place inside the triaxial cell (and with the vacuum pump still connected into the specimen holder), the ram is advanced by the loading system in order to bring it close to contact. The phase of driving up the ram before contact allows a verification of the ram’s friction with its seal to be made. In order to get as close to contact as possible, the x-ray scanner is turned on and the source and detector are translated downwards in order to have the bottom of the specimen in the middle of the image. The detector is then used to create a stream of acquisitions – a specimen radiograph very close to contact is shown in Figure 3.9. This stream of images allows the progress of the ram to be followed very carefully. When the ram is very close to contact, the LVDT is physically set to the beginning of its range (since from this point on the specimen will only be compressed), locked in place and zeroed.

With the axial loading system ready, the confining pressure needs to be added. To do this, the cell is filled with tap water at room temperature. To detect possible leaks in the membrane



Figure 3.9: Radiograph from test COEA01 showing (from top to bottom) the lower porous stone, bottom platen, and the axial loading ram

from this early stage, the filling of the cell is imaged using x-rays. A tube is connected to the tap in the bottom of the cell, through the chicanes in the cabin and to the outside of the cabin, so that the progress of water filling can be controlled with the cabin sealed. A small set of these images is shown in Figure 3.10. If there is a leak, a careful comparison of the images can show water moving on the inside of the specimen. If this is noticed, the specimen is abandoned.




Figure 3.10: Five radiographs of the cell being filled with water, with specimen COEA01 in place

At this point, the specimen is ready for isotropic compression stage. The compressed air coming from a compressor is connected through a pressure regulator, through the chicanes of the tomograph, through a pressure-meter, and finally into the top of cell in order to pressurise the confining fluid. Before applying pressure on the confining fluid, the data acquisition (axial force, axial displacement and cell pressure) is started, with recording every 5 seconds.

The cell pressure is gradually increased from 0 kPa to around 60 kPa over about 20 seconds, whereupon the vacuum pump is stopped, and the vacuum inside the specimen slowly released as the cell pressure is increased to the final desired value (100 or 300 kPa typically).

It is worth remembering that the specimen holder is hollow (to allow suction to be applied to the inside of the specimen). This means that if the membrane is pierced, the confining fluid will flow through the hole, into the specimen, and out of the top of the specimen holder.

At this stage the entire system is checked for water leaks, and if none are found, a first tomographic x-ray scan (using the settings detailed in chapter 2) of the specimen is performed in what is labelled as state “01”, under isotropic compression. During the x-ray scan, the force, displacement and confining pressure data acquisition is slowed to once every 60 seconds. Data acquisition is slowed since the duration of the tests performed in this work (several days) means that record-

ing file can rapidly become large, and no rapid changes in the measurements is expected during a scan.

3.5.2 Triaxial testing

When the first scan has been acquired and checked, the acquisition rate for the axial force, axial displacement and cell pressure is put back to 5 seconds. The ram is now driven up by the loading system at 21 $\mu\text{m}/\text{min}$, which represents a strain rate of just under 0.1% per minute for a specimen measuring 22 mm height. The axial force is closely monitored for the point at which it starts to climb, indicating that the ram is in contact with the specimen, and is starting to load it. At the point where the axial force starts to climb, the axial displacement is recorded and noted as the real zero displacement for this test. Hereafter, loading is continued at 21 $\mu\text{m}/\text{min}$ until a decision is made to interrupt loading.

Often in the tests presented in this work, steps of approximately 1% axial shortening are aimed for, but in cases where the macroscopic response suggests significant changes have occurred in the specimen, scans are performed. Furthermore, some scans for specific purposes are also made – for example some scans very close to each other are performed in test COEA01. When loading, radiographs of the specimen deforming are acquired to be able to identify any sudden events that might occur. These are acquired under the same conditions as the tomographic scans, except that the rotation stage is kept stationary and that no images are averaged for each projection.

Loading is interrupted, when a scan is desired, by simply stopping the loading system's motor; this causes the system to stop moving up. Therefore, when loading is interrupted the axial displacement is simply held constant.

Axial stress relaxations have been noticed to start a few seconds after loading is interrupted, and to continue for a few minutes. The data acquisition rate is maintained at one set of readings every five seconds, to monitor the axial stress relaxation. After approximately 10 minutes, when the axial stress relaxation is essentially over, the tomographic x-ray scan is launched and the data acquisition rate set back to once every 60 seconds.

Large axial stress relaxations have been noticed in specimens left under load overnight; the cause of which has been related to the temperature of the room. The cell is made of PMMA, which has a coefficient of thermal expansion of approximately an order of magnitude greater than that of ordinary steel. Therefore for a positive change of temperature (the heating turning on in the morning in winter) the cell will expand vertically more than the loading system-force meter-ram-specimen-holder system, and therefore will reduce the force measured.

Scans launched immediately after the interruption of loading do not appear to be more blurred than those launched after a given time, which possibly is an indication that no granular re-arrangement is taking place and that the relaxation is perhaps a response of the system as a whole.

At the end of a test, the specimen's response as it is unloaded back to an isotropic state of stress is recorded, after which the cell pressure is released the cell drained and dismantled. The mass of the grains making up the specimen is measured. Each specimen and membrane is kept in a sealed and labelled container.

3.5.3 Naming convention

Throughout this work, the following naming convention is adopted:

- A test or specimen is designated by a 6 character label, where:
 - the first two characters designate the material (CO="Caicos Ooid", OU="Ottawa Uncemented", HN="Hostun sand", AB="Atlanta Ballotini")
 - the second two characters designate the researcher in charge of the experiment

- the last two characters are numbers designating the number of the specimen being tested
- A mechanical state is designated by a series number appended to the label; *e.g.*, HNEA01-01 indicates the first imaged state in a test (HNEA01-00 is reserved for data pertaining to the setup of the test, such as the radiograph of the water filling the cell, or ram getting close to contact)
- An increment, or a transition between two states is indicated by the label, followed by both states; *e.g.*, HNEA01-01-02 indicates the first increment in which deviatoric load is applied

3.6 Experimental campaign

The experimental campaign consists of a study of three different materials with considerably different grain shapes, ranging from the very angular Hostun sand, to the rounded Caicos Ooid. These materials have been tested at different confining pressures, and each test has at least one backup test. The backup test for the Caicos material at high pressure has been selected from previous work (and is labelled COCV02). This section describes the specimens of each type of material and the details of their mechanical responses are discussed.

In total 20 tests have been performed in this study including supplementary backup tests and tests on other materials including glass ballotini. However only the the tests outlined in the following are discussed in detail.

Table 3.2 outlines the tests that will be analysed in this work.

Confining Pressure	Hostun	Caicos Ooid	Ottawa sand	Other Materials
100kPa	HNEA01 HNEA03	COEA01 COEA03 COEA04	OUEA04 OUEA06	ABEA02
300 or 400 kPa	HNEA02 HNEA04	COEA02 COCV02	OUEA01 OUEA02 OUEA03	

Table 3.2: Table outlining the names (6-character labels) of the tests that make up the experimental part of this work, showing what combination of material and pressure each label represents. Bold letters refer to tests that have been analysed in this work

This section will show the macroscopic stress and strain measurements made for each specimen, as well as measurements pertaining to each specimen’s initial state. The presentation is organised into materials and then into the high and low confining pressure cases. The macroscopic response of each individual test is specifically commented in this section, however further analysis will take place in the Chapter 6, where the macroscopic measurements will be supplemented by the micro-scale analysis made possible with the tools developed in this work.

The stress ratio shown (for historical reasons, and because it allows triaxial tests at different cell pressures to be compared easily) is calculated by the axial stress measured (σ_1) divided by the cell pressure imposed (σ_3). The axial stress is calculated by the axial force measured over the cross-section of the specimen in the initial state. The axial force comes from account the raw measurement of force with the contribution of the cell pressure being exerted on the cross section of the ram removed. The membrane correction described in Bishop and Henkel (1962) has not been applied to these data.

The normalised axial shortening is simply the shortening applied by axial compression of the specimen, with respect to its initial height. Axial shortening is typically known as axial strain,

however it is not called as such in this work because when strain localises this quantity is no longer representative of the strain field in the specimen.

Since specimens are tested dry, volumetric strain cannot be obtained from the volume of the pore fluid entering or leaving the specimen. The bulk volume of the specimen is instead calculated from the 3D images obtained throughout the tests; these are shown as circles on the volumetric strain plots (and should correspond to points where axial stress relaxation is visible in the stress response). The bulk volume calculation from the 3D images coming from tomography is detailed in Section 4.4.

General remarks

The dimensions of the specimens created in this work have a certain degree of variability, although the diameter varies relatively little – the most extreme cases being 10.28 mm and 10.92 mm, *i.e.*, approximately two grain diameters. The height of the specimens however is much more variable, in this case the extremes are 22.17 mm and 26.94 mm. This variability is due to the fact that during the preparation of the specimen (as illustrated in Figure 3.6), the top surface of the specimen being created is still inside the mould, and judgement of the height of the specimen can be difficult.

Tables 3.3, 3.4 and 3.5 for each material give the initial conditions of the specimens tested. Measurements of specimen diameter and height, as well as the angles of offset of the specimen from vertical, and the offset of the bottom platen from horizontal are made directly on the 3D images coming from tomography of the initial state (*i.e.*, when the sample is isotropically compressed with the final value of the cell pressure). Slices through the 3D volumes in the initial condition are also given in order to illustrate the tilt of the specimen and platen.

Measurements of global porosity and relative density are calculated using both the specimen's bulk volume (which comes from tomography) and the specimen's solid mass.

The dilatancy angle is calculated from the angle between the volumetric strain rate $\dot{\epsilon}_v$ and the shear strain rate $\dot{\epsilon}_s$, which is in turn calculated as follows: $\dot{\epsilon}_s = \frac{2}{3}(\dot{\epsilon}_{axial} - \dot{\epsilon}_{radial})$.

3.6.1 Hostun sand

Table 3.3 gives details of the Hostun specimens tested (two at 100 kPa confining pressure and two at 300 kPa confinement). Measurements coming from the 3D images acquired in the initial condition "01" of each test are taken with the confinement applied. Relative densities are calculated with respect to the maximum and minimum density reported in Combe (1998) for RF sand which gives a γ_{max} of 16.191 kN/m³ and a γ_{min} of 13.489 kN/m³. The relative densities achieved in the four specimens studied in this work is quite similar (they are dense) however due to the small size of the specimens, and the difficulty in preparation, there is a considerable amount of variability in the relative densities of the specimens, when compared to the variabilities obtained in full size specimens. Figures 3.11 and 3.12 show the specimens' macroscopic axial shortening against stress ratio as well as volumetric strain plots for 100 kPa and 300 kPa confining pressure, respectively.

The two tests performed on Hostun sand at 100 kPa (HNEA01 and HNEA03), have relatively similar macroscopic responses. Both of these specimens start from relatively similar initial conditions with HNEA01 having 83.2% relative density and HNEA03 having 73.1%. HNEA03, however, is considerably longer than HNEA01, and has a considerable tilt to its bottom platen. Both specimens behave as expected for a dense granular material: there is a peak in the specimen's axial stress response, followed by strain softening until a plateau of residual stress is reached. The peak is very close in both cases and corresponds to a stress ratio of about 6.15, which is reached at around 5% axial strain (or shortening). This peak stress ratio corresponds to an axial stress of $6.15 \times 100kPa = 615kPa$ axial stress (*i.e.*, σ_1 in this case). This corresponds to a friction angle of 46.1°, which corresponds well to the value reported in Combe (1998) for

Hostun S28 sand of 46.2° at 100 kPa. The representativity of these small specimens is discussed later, in Section 3.7.

Leading up to peak, both specimens are quite stiff in their initial loading, and HNEA01's more perfect geometry is likely to be the reason behind the specimen's stiffer initial response and is also likely to be one of the reasons behind its slower progression towards a residual stress state when compared to HNEA03.

A residual stress ratio of 4.15 ($\phi' = 37.7^\circ$) is reached earlier by specimen HNEA03 at around 13.0% shortening – specimen HNEA01 resists longer, reaching a higher residual stress ratio of 4.4 at around 13.3% axial shortening. In both cases the residual stress at the end of the test is confirmed by the volumetric response of the specimen, which reaches a state of zero dilatancy by the end of the experiment.

The volumetric response for both specimens is dilatant with approximately the same dilatancy angle at peak stress ratio (17.0° and 18.1° for HNEA01 and HNEA03, respectively). Specimen HNEA03 can be seen to have considerably more contraction than specimen HNEA01, possibly due to its lower porosity. Specimen HNEA01's dilatancy angle appears to reduce around 10% axial strain, which goes some way to explain this specimen's slower approach to residual stress; it could be hypothesised that the specimen is more resistant to getting to critical state. In specimen HNEA03's volumetric response there is a very anomalous point coming from the scan of state HNEA03-06, which may reflect the error bounds of the volume measurement technique.

The two specimens tested at 300 kPa (HNEA02 and HNEA04) are quite similar, both starting from high relative densities (95.4% and 84.3% respectively). Geometrically speaking, they have approximately the same deviations from a perfect cylinder. However, in their axial stress response it is clear that there is an offset in the peak value of stress ratio measured. HNEA02 has a stress ratio at peak of 5.5 ($\phi' = 43.8^\circ$) whereas HNEA04 peak stress ratio is 5.35 ($\phi' = 43.2^\circ$), although this ratio is reached in both cases around 7.2% shortening. The peak stress reported for Hostun S28 in Combe (1998) is 43.3° , closer to the lower value of HNEA04.

The specimens reach peak stress with an increased initial stiffness compared to the test at 100 kPa confinement (stiffness cannot be directly measured as a gradient on this plot), and from their stress response do not appear to reach a residual state in the range shown; this is confirmed by the volumetric responses, which show a reduction in dilatancy angle, but without it reaching zero by the end of range shown. HNEA04 reaches a residual stress ratio of 4.0 ($\phi' = 36.7^\circ$) at around 15.3% shortening (*i.e.*, off the graph). Qualitatively speaking it appears as though specimen HNEA02 will have a higher residual stress than HNEA04.

As for the tests at 100 kPa on this material, the specimen with a lower relative density (HNEA04) has more initial contraction, but both specimens achieve similar angles of dilatancy (15.0° and 14.8° for HNEA02 and HNEA04 respectively).

It can be noticed that in the increment HNEA02-10-11 (*i.e.*, between 12.4% and 13.9% shortening) the axial stress readings have a different appearance than the other increments – this is down to an operator error of forgetting to increment the reading rate of the measurement after the end of a scan. The data are thus considerably more sparse in this increment. These measurements are summarised in table 3.6.

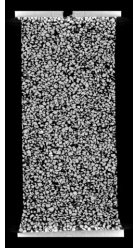
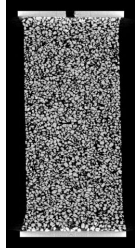
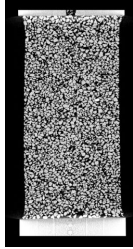

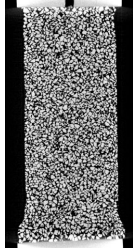
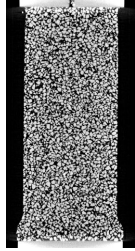
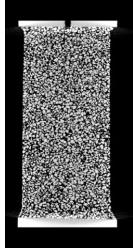
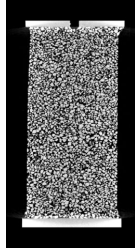
Test name	Confining pressure	Diameter	Height	Mass	Porosity (D_R)	Maximum tilt of specimen	Tilt of bottom platen	Slice showing platen tilt	\perp Slice
HNEA01	100 kPa	10.70 mm	24.31 mm	3.55 g	39.7% (83.2%)	0.8°	1.9°		
HNEA02	300 kPa	10.76 mm	23.11 mm	3.36 g	38.2% (95.4%)	0.5°	1.2°		
HNEA03	100 kPa	10.86 mm	26.94 mm	3.85 g	40.9% (73.1%)	0.5°	2.8°		
HNEA04	300 kPa	10.92 mm	23.46 mm	3.42 g	39.6% (84.3%)	0.3°	1.0°		

Table 3.3: Table showing initial setup for each specimen of Hostun sand. Measurements of diameter, height, porosity and tilt angles all derive from the x-ray tomography images




Figure 3.11: Normalised Axial Shortening plotted against Stress Ratio (top) and Volumetric Strain (bottom) for Hostun specimens at 100kPa




Figure 3.12: Normalised Axial Shortening plotted against Stress Ratio (top) and Volumetric Strain (bottom) for Hostun specimens at 300kPa

3.6.2 Ottawa sand

Table 3.4 describes the initial conditions of the Ottawa specimens tested in this work. Measurements coming from the 3D images acquired in the initial condition "01" of each test are acquired with the cell pressure applied. The maximum and minimum relative density values are taken from Lee *et al.* (2007), who report a specific gravity of Ottawa sand grains as 2.65 and minimum void ratio of 0.50, which gives a maximum density γ_{max} of 17.33 kN/m³. The maximum void ratio is given as 0.85, which gives a maximum density of 14.05 kN/m³.

Figures 3.13 and 3.14 show the macroscopic axial shortening *vs.* stress ratio and volumetric strain responses of the Ottawa sand specimens tested in this work at 100 kPa and 300 kPa confinement respectively.

Specimens OUEA04 and OUEA06 were tested at 100 kPa confining pressure, and their stress-strain responses can be seen to be very similar. Both specimens have similar initial stiffness, OUEA04 reaches a peak stress ratio of 5.50 ($\phi' = 43.8^\circ$) at around 4.7% axial shortening, whereas OUEA06 has a peak of 5.66 ($\phi' = 44.4^\circ$) at around 4.3% shortening. Both specimens undergo strain softening until reaching a plateau in the stress ratio of 3.75 ($\phi' = 36.0^\circ$) at around 13% axial shortening. The apparent plateau at 13% axial shortening cannot be confirmed as a residual stress state from the volumetric response for specimen OUEA04, since the angle of dilatancy does not reach zero by the last scan. Test OUEA06 does however reach a state of zero dilatancy angle, thus confirming that it has reached the residual stress state. The dilatancy angles at peak are 16.1° and 20.2° for OUEA04 and OUEA06, respectively. Despite very little difference in initial porosity, specimen OUEA06 seems to contract a little less than specimen OUEA04.

There are some stress data missing for the increment OUEA06-07-08 (between 5.76% and 7.80% axial shortening) because the LVDT that measures the displacement imposed on the specimen went out of range. The straight line visible on the graph simply connects the first and last recorded points and is left as is in order to highlight this deficiency. Furthermore, specimen OUEA06 suffers from a small invasion of water during the test, but the volumes of water coming in are small, and so the stress applied by the confining pressure is thought not to be lost.

As for the Hostun specimens, the two Ottawa sand specimens tested at 300 kPa confinement have a slight offset in the measured stress response, although the curves have qualitatively the same shape. Both specimens reach their peak stresses at approximately 5.2% axial shortening, although OUEA02's peak stress ratio (5.20 *i.e.*, $\phi' = 42.6^\circ$) is higher than OUEA03's (5.02 *i.e.*, $\phi' = 41.9^\circ$). Both specimens appear to tend towards similar values of residual stress ratio, achieving them around 11.6% axial shortening (approximately 3.7 for OUEA02 and 3.5 for OUEA03). This is confirmed, by the dilatancy angle in the volumetric response which reaches 0 by the end of both tests.

As seen in the Hostun specimens, both specimens have a similar dilatancy angle at peak (16.1° and 17.1°), but the less dense of the two specimens has a more significant phase of contraction. Another trend that is confirmed is that the initial stiffness of the specimens tested at high pressure is higher than that at lower pressure.

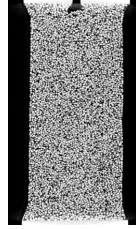
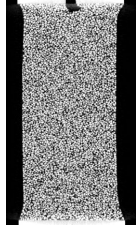
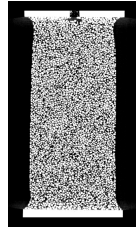
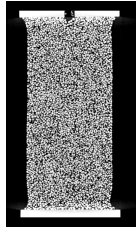
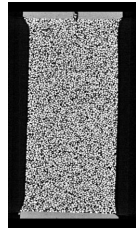
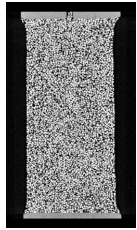
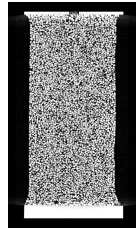
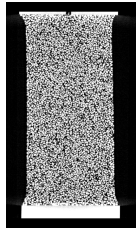
Test name	Confining pressure	Diameter	Height	Mass	Porosity (D_R)	Maximum tilt of specimen	Tilt of bottom platen	Slice showing platen tilt	\perp Slice
OUEA02	300 kPa	10.71 mm	24.65 mm	3.77 g	35.6% (84.9%)	0.3°	0.1°		
OUEA03	300 kPa	10.28 mm	23.45 mm	3.34 g	35.3% (87.0%)	1.3°	0.7°		
OUEA04	100 kPa	10.76 mm	23.46 mm	3.75 g	31.4% (111.7%)	1.3°	1.1°		
OUEA06	100 kPa	10.55 mm	23.02 mm	3.58 g	31.4% (112.3%)	0.4°	0.5°		

Table 3.4: Table showing initial setup for each specimen of Ottawa sand. Measurements of diameter, height, porosity and tilt angles all derive from the x-ray tomography images




Figure 3.13: Normalised Axial Shortening plotted against Stress Ratio (top) and Volumetric Strain (bottom) for Ottawa specimens at 100kPa




Figure 3.14: Normalised Axial Shortening plotted against Stress Ratio (top) and Volumetric Strain (bottom) for Ottawa specimens at 300kPa

3.6.3 Caicos ooids

Table 3.5 gives details of the Caicos ooids specimens tested (three at 100 kPa confining pressure and two at 300 kPa confinement, one of which was done prior to this work by Cino Viggiani and Pierre Bésuelle). Measurements coming from the 3D images acquired in the initial condition "01" of each test are acquired with the confinement applied. No relative densities are reported since this material has not been systematically studied, however, based on the densities reported for the Hostun and Ottawa materials, the densities achieved here are expected to be relatively high.

Specimen COEA01 has been prepared by dry pluviation with the system used for the other materials, whereas the other specimens have been prepared with a much smaller drop height. Interestingly the other specimens appear to have a higher density than COEA01.

Figures 3.15 and 3.16 show the specimens' macroscopic axial shortening against stress ratio as well as volumetric strain plots for 100 kPa and 300 kPa confining pressure respectively.

Three specimens were tested at 100 kPa since the backup test COEA03 did not confirm satisfactorily the response seen on COEA01. As with the other materials tested, these dense specimens show a stress peak, followed by strain softening. In all three tests a plateau in the stress ratio of 3.70 is reached around 11% axial shortening for tests COEA01 and COEA04, and is not quite reached by specimen COEA03 by the end of its test.

Peak stresses in a carbonate sand are expected to be high as stated in Coop (1990), however the peak stress ratio of COEA01 is very large, 6.79, which means a friction angle of 48° that appears at an axial shortening of 3.64%. This is not confirmed by the other two backup tests, which appear to have their peak stress ratio a little earlier (3.4% axial shortening) with a ratio of 5.77 giving a friction angle of 44.8° . Since the residual stresses are the same for these tests, the very high peak stress cannot be ascribed to an error in the instrumentation or in the data treatment. Furthermore, even the two backup tests have surprisingly high friction angles for a rounded sand.

The initial stiffness of COEA03 is considerably higher than the other two tests, and its post-peak strain softening is more gradual than for the other two tests.

Despite variations in initial porosity similar to the other specimens of the other materials, in this case the volumetric response is very similar for all specimens, with a dilatancy angle at peak of 20.4° for COEA01 and COEA04 and an angle of 19.3° for COEA03.

The reasons for the differences between these tests – principally between COEA01 and the other two backup tests will be a challenge for the micro-level analysis made with the tools proposed in this work.

The two tests at higher cell pressure (COEA02 at 300 kPa and COCV02 at 400 kPa) again have a higher initial stiffness than their lower cell pressure counterparts – with COCV02 having a lower initial stiffness than COEA02. Both specimens reach their peak in stress ratio at around 5.2% axial shortening with COEA02 having a value of 5.97 ($\phi' = 44.5^\circ$) and COCV02 having one of 5.85 ($\phi' = 45.1^\circ$). Both specimens seem to reach considerably different residual stress ratios after 11% axial shortening: COEA02 reaches a value of 3.65, whereas COCV02 gets to 3.11.

Immediately upon reloading after a scan, some additional resistance to shear seems to be created by a mechanism that is not fully understood; at some points of these two tests this adds a visible step above the loading curve upon reloading. After this step, the specimen regains the response curve that it was following before the scan. These phenomena are visible to some extent in the tests on Caicos ooids at 100 kPa, to a lesser extent in the Ottawa tests and not at all in the tests on Hostun sand. It has also been confirmed in with on tests on glass ballotini at 100 kPa (which are not shown here) this same apparatus, so a possible hypothesis is that this is linked to grain roundness, possibly some contact rearrangement. This phenomenon, that acts over small strain increments, and appears to be linked with time (the length of the scan), is not easy to study by x-ray tomography and the current experimental setup. Qualitatively speaking

work on viscous effects in granular materials (Enomoto *et al.*, 2009) shows some very similar trends, which are subjects for future research.

The volumetric response of COEA02 shows some contraction, followed by dilation, with a dilation angle of 16.3° , which reduces to 0° as the specimen reaches its residual stress state. COCV02 has many fewer scans during deformation, which means that the data for the volumetric strain are sparse. By eye, the dilatancy angle at peak looks comparable to that of specimen COEA02.

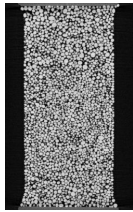
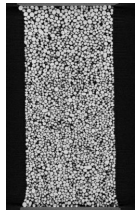
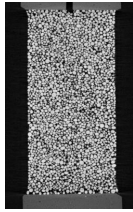
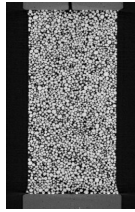
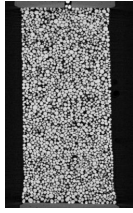
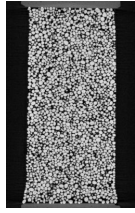
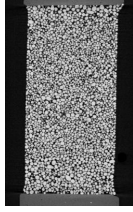
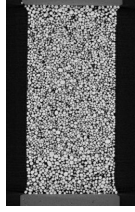
Test name	Confining pressure	Diameter	Height	Mass	Porosity (D_R)	Maximum tilt of specimen	Tilt of bottom platen	Slice showing platen tilt	\perp Slice
COEA01	100 kPa	10.74 mm	24.16 mm	3.85 g	35.2%	0.0°	0.6°		
COEA02	300 kPa	10.55 mm	22.17 mm	3.55 g	33.2%	0.0°	0.9°		
COEA03	100 kPa	10.80 mm	23.44 mm	3.84 g	32.1%	1.1°	1.0°		
COEA04	100 kPa	10.80 mm	22.73 mm	3.76 g	31.9%	1.2°	0.5°		
COCV02 (Exxon_04)	400 kPa	10.86 mm	21.74 mm	3.53 g	33.8%	2.2°	1.3°		

Table 3.5: Table showing initial setup for each specimen of Caicos Ooids. Measurements of diameter, height, porosity and tilt angles all derive from the x-ray tomography images




Figure 3.15: Normalised Axial Shortening plotted against Stress Ratio (top) and Volumetric Strain (bottom) for Caicos specimens at 100kPa




Figure 3.16: Normalised Axial Shortening plotted against Stress Ratio (top) and Volumetric Strain (bottom) for Caicos specimens at 300kPa

3.7 Representativity of specimens and their responses

Mechanical pertinence of the tests and representativity of material responses is clearly a concern with the use of such small specimens. This section discusses this issue with respect to Hostun sand, since several good references for triaxial tests are available for Hostun sand. A number of very complete studies of the behaviour of Hostun sand under triaxial compression exist (*e.g.*, Bousquet *et al.*, 1993), but the results of Combe (1998) are used herein, since their raw data are still easily accessible. In both the sources provided above, triaxial tests on specimens with a slenderness ratio of 1 (diameter = height) are tested. If the small specimens of Hostun sand tested in this work are globally comparable to much larger standard triaxial tests the same is expected to hold true for the small specimens of the other materials tested.

In order to get a sense of scale, the specimens tested in this work have a mass between 3 and 4 grammes, whereas a dense specimens of Hostun sand of “standard” size have a mass of approximately 2.4 kilogrammes (depending on density of course). The very different sizes of the specimens may have an effect on the mechanical response measured, since with the small specimens it almost makes more sense to measure the specimen’s dimensions in grains rather than millimetres: for example the specimens tested in this work have diameters of around 10.5 mm which means, in the case of Hostun sand (where the D_{50} is 388 μm), that the diameter of a specimen is $31 \times D_{50}$.

When physical specimens start to get small with respect to the grain size, it is possible that some deformation mechanisms simply are not possible due to the geometry and size of the specimen: for example the complicated patterns of shear band observed by x-ray tomography in Desrues *et al.* (1996) simply are not possible in the specimens tested in this work due to the number of grains involved.

Anecdotal evidence provided by some post-mortem photographs in Combe (1998) shows that at least some specimens failed with a single shear band. This is the same failure mode with which the small specimens tested in this work deform. The conclusion to be drawn is that despite the small number of grains making up the small specimens tested – which has been hypothesised to limit the material’s freedom of choice in deformation mechanisms – the material under these conditions of macroscopic stress and strain is able to deform in the same way for both sizes of specimen.

Figure 3.17 shows HNEA01 compared to a single triaxial compression test (TS28D01) performed at $\sigma_3 = 100$ kPa on a specimen measuring 100.2 mm height and 99.9 mm diameter. In the same graph, HNEA02 is compared to a test with the same initial effective mean stress, TS28D03 measuring 100.1 mm height by 99.5 mm diameter.

Both full-size specimens have a less-pronounced peak, which occurs later in the test. The value of stress ratio at the peak is much more reproducible: at 100 kPa confinement the peak is at 6.15 for HNEA01 and 6.16 for TS28D01; at 300 kPa confinement the peak is at 5.50 for HNEA02 and 5.46 for TS28D03. The initial stiffness of the small specimen at 100 kPa is considerably greater than that measured for the large specimen – at 300 kPa the initial stiffnesses of the two types of specimen coincide very closely. The angle of dilatancy of all the specimens at peak stress ratio is comparable.

The small specimens tested reach their respective peak stresses and residual states much more rapidly than their larger-sized equivalents. The changes observed in the material response when changing confinement from 100 kPa to 300 kPa however appear in both sizes of specimen: at higher confinement the peak stress appears later and the friction angle observed at the peak is lower. The fact that the small specimens reach peak and residual stress states faster might be an indication that at the micro-scale the deformation mechanism that leads to failure – *i.e.*, the development of a single shear band – develops more readily in the small specimens when compared to the larger ones. This in turn could be for a number of different reasons:

- Due to the technical difficulty in the creation of such small specimens, their geometry is very




Figure 3.17: Graph showing axial shortening *vs.* stress ratio and volumetric strain for two tests on Hostun sand coming from this work, and two test from Combe (1998) for comparison

likely less perfect than larger-sized specimens. This in turn means that more significant inhomogeneities arise and potential locations for shear banding in the specimen may be activated earlier in the specimen’s loading. Furthermore the same heterogeneities are likely to have a more significant effect on small specimens.

- The larger slenderness ratio of the small specimens may also favour shear banding since a shear band’s inclination in triaxial compression is larger than 45° from horizontal, it will be more geometrically difficult for a single shear band to appear in a specimen of slenderness ratio 1, than in a longer specimen where it can develop without interaction with both ends of the specimen.

The conclusion to be drawn is that there are some clear size effects in the macroscopic responses of the specimens of different dimensions that have been compared. The peak stresses and volumetric responses are however close in both cases, and the changes between the two confinement values are coherent between the small and the large specimens. The small specimens can therefore be considered to be mechanically representative of the larger specimens tested in previous work.

3.8 Summary of mechanical responses

Table 3.6 shows an overview of the specimen responses obtained from the experimental campaign. Values have been weighed according to the quality and confidence in each test.

Material	Confining pressure	Peak stress ratio σ_1/σ_3	Residual stress ratio	Angle of dilatancy at peak
Hostun sand	100 kPa	6.15 ($\phi'_p = 46.0^\circ$)	4.15 ($\phi'_r = 37.7^\circ$)	17.5°
Ottawa sand	100 kPa	5.55 ($\phi'_p = 44.0^\circ$)	3.75 ($\phi'_r = 36.0^\circ$)	18°
Caicos ooid	100 kPa	5.77 ($\phi'_p = 44.8^\circ$)	3.70 ($\phi'_r = 35.0^\circ$)	20°
Hostun sand	300 kPa	5.4 ($\phi'_p = 43.4^\circ$)	4.15 ($\phi'_r = 37.7^\circ$)	14.9°
Ottawa sand	300 kPa	5.1 ($\phi'_p = 42.2^\circ$)	3.70 ($\phi'_r = 36.0^\circ$)	16.5°
Caicos ooid	300 kPa	5.77($\phi'_p = 45.2^\circ$)	3.65 ($\phi'_r = 34.4^\circ$)	16.3°

Table 3.6: Table showing averaged highlights of the mechanical responses (weighted by confidence) measured in tests presented in this work. COEA01 is not included in these data

Chapter 4

Image analysis

This chapter details the analysis of the 3D images coming from the triaxial tests performed *in-situ* in the Laboratoire 3SR x-ray scanner. These 3D images *are* the micro-scale measurements for the experiments described in Chapter 3, and so all the micro-mechanical analysis in this work is done starting from these images, ultimately in preparation for the kinematic analysis in Chapter 5. These 3D images represent a field of a reconstructed x-ray attenuation value, evaluated all over a specimen in a given state and are therefore greyscale (*i.e.*, they are a field of a single, continuously variable, scalar). All the specimens in the campaign have been imaged in several different states of increasing axial strain. In these images grains are directly visible to the naked eye thanks to two facts: first, the grains' x-ray attenuation is much higher than the air that surrounds them (*i.e.*, the contrast is good between grain and void) and second, the zoom is such that there are typically between ten and thirty pixels across a grain's widest diameter, allowing them to be easily distinguished in the images. Individually distinguishable grains are visible in Figure 4.1, which shows sections, sliced horizontally through the 3D volume of two different samples. It is very important to note that although it might appear from these reconstructed slices that there is quite a wide grain size distribution, this impression comes from the fact that the images presented are measurements of x-ray attenuation averaged over a unit thickness of only one voxel (*i.e.*, 15 μm). This slice will have cut some grains through their thickest portion and others in less-thick portions, making them appear as larger or smaller (respectively) in this 2D slice. Another visual artefact from the process of extracting slices from a 3D image is that some grains seem not to touch any other grains *on this slice*; they will however necessarily touch some other grains on other slices.

This work follows the philosophy of discrete (*i.e.*, grain-based) measurements where possible. However, when a spatially-varying field such as porosity (which has a strict and familiar geotechnical definition) is under consideration, it will be evaluated over a regular grid (as is natural since 3D images provide a convenient grid). A field like porosity might also be measured with respect to the underlying discrete granular skeleton which could be used as a reference system, the field being evaluated for each grain. Since this is purely a matter of philosophy, where practical and where it is deemed interesting, comparisons of fields evaluated on a grid, or on the granular skeleton are made. The evaluation of a field on a regular grid will be called a "continuum" analysis, since the choice to ignore the underlying granular structure is to treat the material as a continuum, where it is natural to specimen the material points at a regular interval.

Layout of chapter

Since the overarching objective of this work is to analyse the mechanical behaviour of sand at the grain scale, the grains that compose the specimens under study need to be individually accessible and measurable. This means that grains in some way need to be defined, reliably, in each 3D




Figure 4.1: Horizontal slices taken from initial states of (a) HNEA03 and (b) COEA03. The grains (light against a dark background) are clearly visible in these images. A small shell as well as some very porous grains are visible in slice (b), the former reminding the reader of the marine origin of this sand. In these images white represents high x-ray attenuation coefficients and black low x-ray attenuation

image (which only contains information about x-ray attenuation). Section 4.1 describes how grains, and the contacts between them, are defined in the 3D images with a sequence of image analysis operations. Once grains are defined in a 3D image, each grain can be extracted and various measurements can be made on the voxels that make it up (such as measurement of volume, centre of mass, and maximum and minimum dimensions). Contacts coming from image analysis are in turn also labelled, so that the labels of the grains they connect is recorded. Contacts are also individually characterised through a set of measurements. Once a set of measurements exists for each grain (or contact), these can be processed into different static measurements (*i.e.*, made on a single image) of the granular skeleton, such as coordination number or grain size distribution. The measurements of grain and contact properties are detailed and characterised in Section 4.2. An unconventional implementation of the conventional “continuum” tools for the measurement of porosity (or void ratio) are then described in Section 4.3. To finish the analysis possible on a single 3D image, a simple technique for measuring the total volume of a specimen in a given configuration is illustrated in Section 4.4.

The chapter that follows (Chapter 5) will take advantage of many of the results of this section, in order to follow and characterise the deformation of granular media in 3D.

4.1 Processing of individual 3D images for the identification of grains and contacts

This section details how individual grains are defined and measured in each greyscale 3D image coming from tomography. This is an essential step for analysis at the “micro”-scale, *i.e.*, the scale of the grain for sand. First, the procedure for defining the solid and void phases in the images (based on a threshold value of x-ray attenuation) is described in Section 4.1.1. The solid phase is then split into separate grains using a marker-controlled watershed algorithm (Section 4.1.2). Section 4.1.3 outlines how the voxels belonging to each grain are then assigned a label (unique positive integer) so that they can be extracted from a labelled 3D image. This allows different

voxel-based measurements to be made on each grain in the following section.

4.1.1 Binarisation

For the two-phase (air and grains) media under study, it is essential to be able to precisely distinguish the solid (grain) phase from the void, to be able to define (after further treatment) individual grains. Binarisation describes the process of creating a binary (*e.g.*, black-and-white) image where all solid phase voxels contain one value, and all other voxels contain another value.

The greyscale images coming tomography are not already binary, despite being images of a two-phase medium, for three principal reasons: first, although the pore fluid is expected to have a homogeneous distribution of x-ray attenuation, the same cannot be said for the grains being imaged, which may present some character on their inside. The second reason for the range of greyvalues in the images coming from tomography is simply noise – there is a certain amount of noise generated in the acquisition of x-ray images, and possibly more added in the 3D reconstruction of data. Since noise is expected to be relatively random, this will have the tendency of generating different shades of grey. The third cause of the range of greyscale values in these data is the partial volume effect, described below, which occurs at the interface between phases.

Since there is a range of greyvalues in the images being treated, a threshold value must be chosen in order to separate the two phases. The way in which this is chosen is also detailed below.

Partial Volume Effect

Digital images are composed of discrete units (pixels in 2D and voxels in 3D). In any digital image acquired with an array of sensors, each pixel value will be a measurement integrated over the small area of the sensor. If a two-phase medium is being imaged, the values of pixels well inside either phase will correspond to the measurement of that phase alone (assuming no noise). Any pixels recording the interface between phases will have a value that results from the integration of measurements of both phases over the small area of the pixel – the result being an intermediate value. This is typically known as the “partial volume effect”. A pixelated image of two phases will chiefly contain two different values (the measurements corresponding to the two phases) plus a certain amount of measurements in between coming from the partial volume effect. A histogram of the measurements in this image would show two distinct peaks with some information between them. threshold

In the 3D greyscale images used in this work, the partial volume effect accounts for a region of about 3 pixels width.

Choice of threshold

Maps of the solid phase of the studied materials are constructed by binarising the reconstructed 3D images of x-ray attenuation. The radiograms acquired during a tomographic scan are subject to the partial volume effect as well as noise, and consequently both these effect are manifested in the reconstructed images. Since the solid phase is known to have a higher x-ray attenuation value than air or water (the confining fluid), if a unique threshold value of x-ray attenuation (above which voxels are assigned to the solid phase) can be accurately chosen, a reliable map of the solid phase can be obtained. Ideally, the chosen threshold value is between the attenuation values of the solid and pore phases, and will select a certain amount of the partial volume voxels with higher values.

In order to choose the threshold value, two options are typically available: to make a choice based on the image (typically on the distributions of greyvalues in the image), such as Otsu’s method (see Otsu, 1975), or to determine a threshold from a physical measurement. For this

work, the latter approach has been chosen chiefly for reliability. The physical measurement which the threshold is chosen with is the volume of the solid phase (other measurements could include macroscopic porosity, for example). The dry mass of a specimen is measured post-test, and, since the densities of the grains are well-known, the solid volume is easily calculated. Knowing the volume in μm^3 that a single voxel represents in a given 3D image, the value of the threshold is chosen such that it includes the number of voxels corresponding to the measured volume of the solid phase. This is illustrated in Figure 4.2 – the threshold value can be seen to lie in between the two peaks in the histogram.




Figure 4.2: Illustration of the application and calculation of a threshold value to a slice from test HNEA03. The voxels selected by the threshold (in red) show that the darker partial volume voxels on the outsides of the grains lie below the threshold and are thus not selected.

If an image contains some voxels which have a high coefficient of x-ray attenuation but which do not represent grains, this will cause a too-low threshold to be calculated for the image. This is typically the case for the “porous” stones used for the tests. In the images acquired these elements are of a density similar to the grains, and so for the calculation of the threshold, these high values are not taken into account (they are manually set to zero in the image, when calculating the threshold).

It has been observed that the threshold values selected always lie in between the two peaks. The threshold is calculated once per test, in the first state and is re-applied to all subsequent scans in the same test, which are acquired under the same conditions.

4.1.2 Separation (or segmentation)

Separation (often also called Segmentation – see the Glossary) refers to the process of separating the single interconnected solid phase into individual particles. The watershed family of algorithms (generally used in this type of problem) has been chosen in this work, and a variety of different implementations have been tested. The quality of the separation step is crucial for any grain-based analysis: if the separated particles do not correspond to the physical reality of grains (*i.e.*, if grains not separated from each other are considered as a single particle, or if a single grain is split into several particles), micro-scale analysis will be impaired. The watershed algorithm is highly iterative, and so can be a serious bottleneck in the treatment of the copious data coming from the entirety of the experimental campaign. For this reason, a highly optimised commercial

code (Visilog[®]) is principally used in this work.

This section details the main steps of operation of a watershed algorithm. The first step is to identify markers, which are seeds of the watershed algorithm placed with certainty inside each grain. Thereafter these seeds are grown according to a segmentation function by the watershed algorithm until they touch, at which point they are separated.

Markers

The watershed implementations used in this work are of the “marker-controlled” family of watersheds, which are a first line of defence against over-segmentation and over-sensitivity to noise (see Beucher, 1991).

Markers are a small volume indicating a zone considered to be definitely inside a single grain. Markers are passed to a watershed algorithm (described later), that is responsible for the separation of particles from each other, starting from these markers. The number of markers that the algorithm is supplied with will be exactly the number of individual particles that come out of the process: if one real grain contains two markers, it will necessarily be split in two particles. Similarly, if a system of two contacting grains has only one marker, this will be identified as a single particle. Therefore, a good selection of markers is the first important ingredient for a good separation that respects the real physical distinction between particles. Markers are typically calculated by a series of morphological operations on a binarised image. Most implementations for this kind of problem start by calculating the Euclidean Distance Map (EDM) on the to be separated (in this case the solid phase); see Figure 4.3 an example of an Euclidean Distance Map calculated on a 2D image. An Euclidean Distance Map is calculated on a binary image, such that for every pixel in the phase of interest, its flat value (which simply denotes which phase it is in) is replaced by its minimum (Euclidean) distance to the other phase. A solid pixel touching a void pixel would have a small value, whereas a pixel inside a non-hollow particle will have a relatively high value. A 3D image of markers is then obtained by identifying and joining (merging) local maxima in the Euclidean Distance Map. The “localness” of the maximum is an option in certain implementations, which is often no more complicated than simply the difference in values that will be added around the absolute maximum (*e.g.*, if the “merge local maxima” option is set to 5, and a maximum is at a value (on the EDM) of 10, then all voxels contiguous with that maximum having a value and greater or equal to $10-5=5$ will be included in the marker). This may result in other local maxima being included in the principal maximum, which is desirable to avoid non-physical over-segmentation with too many markers). This procedure is illustrated on a 2D slice of a test on glass ballotini in Figure 4.3, which offer an ideal test case due to their near spherical shape and, therefore, near circular cross sections and small regions of contact.

For images of circular or spherical particles, the markers-placement technique described above is very successful at placing markers. However, on the 3D images of the natural materials studied in this work, there are some situations where markers are placed incorrectly. These are illustrated in Figure 4.4 (in 2D, for simplicity). The reasoning made in the illustration also works in 3D, but care must be taken, especially with the first case where a flat contact is enough to cause problems in 2D, in 3D a large planar contact would be necessary to have this kind of artefact.

As Figure 4.4 suggests, placing markers perfectly inside each grain is difficult: considering the large number of grains in the specimens (not less than 50,000), some of the situations outlined in Figure 4.4 are likely to occur occasionally. The results of these errors are shown in Figure 4.5 – a) shows the result when there are two markers per grain, and b) shows when a single marker is present when there are two grains. The terms (present in the figure) used to denote these situations are “over-separation” in the case where too much separation occurs (*i.e.*, there are locally more markers than grains), and “under-separation” in the case where there are too few markers compared to the number of grains. If markers are placed in such a way that grains are consistently over- or under-segmented, then it means that the automatic placement of markers needs to be tuned. However in the case of the images treated in this work both over and under-

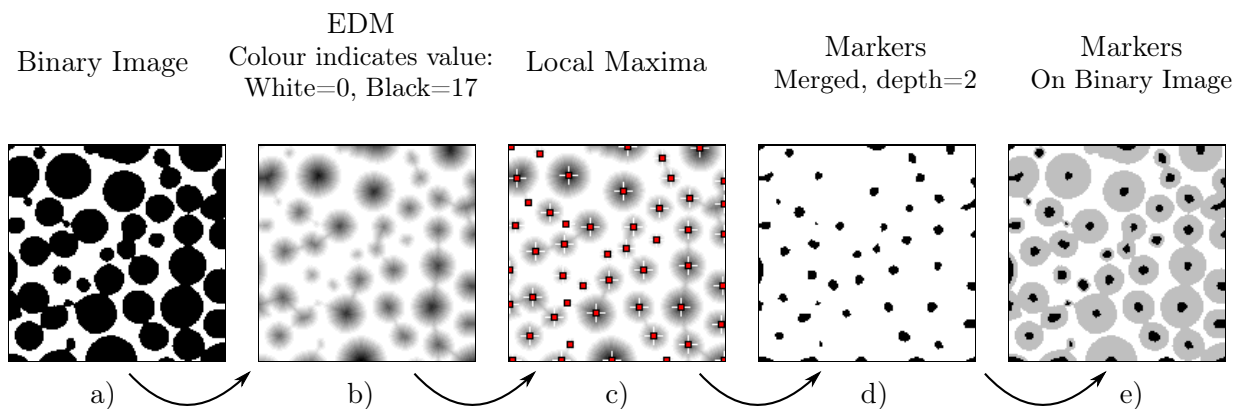


Figure 4.3: Illustration in 2D of the standard procedure for generating markers in an image containing grains. **a)** Shows a binarised detail of a slice from test on glass ballotini (ABEA02). **b)** Shows the Euclidean Distance Map (EDM) for the solid phase (black) of the binary image. **c)** Shows the points which are recognised as local maxima. **d)** Shows the maxima merged to a depth of 2 with the surrounding points. These merged maxima are the markers for the binary image **a)**. **e)** Shows the markers plotted back onto the binary image (faded to light grey). The procedure is successful, since all the grains distinguishable in image **a)** have a marker placed inside them

segmentation occasionally occur. The result of this is that the separated images coming from this procedure are necessarily flawed. The separated images are however deemed (after visual inspection) of a sufficient quality to make measurements.

There is space for improvement in the automatic placement of markers. Due to the large number of images acquired in this work, and also the large number of grains in each image, manual placement of markers is not practically realistic for a whole specimen (but was tempted in the subvolume published in Andò *et al.* (2012b), with good results). Since it is relatively easy to distinguish individual grains by eye in the slices, it is likely that this can be translated into some morphological metrics, which may help in the placement of markers. Some simple marker detection approaches such as those presented in Matsushima *et al.* (2004), include several iterations of a 3D erosion all the unseparated solid phase, with the objective of being left with just the centre of the grains. This however is considered no less vulnerable to error than the general technique described above – with the added risk of deleting some small grains. Better placement of markers is a subject worthy of further attention.

Another, much heavier (from a computational point of view) possibility for the improvement of the placement of markers could be based on the results of the separation process. There are some *a priori* data on the expected morphology of grains (in terms of angularity, volume, and curvature – all of which can be measured on a set of voxels) that could be used to detect badly segmented grains (one grain that has been split, or two grains stuck together). Once detected, a marker-correction algorithm could try to join or split markers by locally modifying some of the parameters used for the initial marker map. This has not been investigated further, since the development of powerful image treatment techniques remains a subject of great interest in computer vision. One could also use the output of the grain tracking described in chapter 5 to also correct the grain separation process. This is discussed briefly in the potential improvements to be made to the grain-tracking procedure called ID-Track in Section 5.3.8. The fact that a lot of information can help improve the placement of markers, is that the whole separation phase of the image treatment is the weak link in the chain of the analysis of these images.

Problematic Cases for the Automatic Placement of Markers




Figure 4.4: Diagram outlining in 2D some situations in which the automatic calculation of markers (as the maxima of an Euclidean Distance Map) used for separation may not give physically correct results. **a)** [Hostun Sand] When a contact is very flat, the euclidean distance to the pore may be sufficiently high to give one maximum of the distance map (and therefore just one marker) for the two grains in contact. In this case the maximum actually lies over the contact! **b)** [Caicos Ooid] If a grain contains even just one voxel of pore on its inside, this deeply distorts the euclidean distance map, with unpredictable results (multiple particles are the usual result). **c)** [Caicos Ooid] With elongated grains, the maximum of the euclidean distance map can be a ridge – and if the grain is slightly thinner in the middle, the ridge may be broken – causing two markers to be placed at the end of the grain)

The watershed algorithm

With a set of markers placed inside each grain (see above), the challenge of separating the solid phase into individual particles remains. Consider a very simple 3D system consisting in two binarised grains (each with a marker correctly placed inside). If the grains themselves are not in contact, no separation is actually necessary. If however, the two grains are in contact, then some subtlety is required in the expansion of the markers, so that the markers meet at the contact point between the two grains. This goes to show that some strategy (namely the watershed algorithm) is necessary for dealing with markers that may come into contact with each other when expanded.

The watershed algorithm was first proposed in Beucher and Lantuejoul (1979), and has been developed with substantial improvements since. The watershed algorithm's operation is inspired by the geographical metaphor of the immersion of a landscape, with water being supplied at the markers (which are typically in the troughs of the landscape). The landscape is progressively




Figure 4.5: Figure showing the two main types of separation error (in Hostun sand). **a)** shows a single grain which is split into two particles and **b)** shows two grains which have not been separated into two distinct particles.

flooded with different markers expanding based on the advance of the water they supply. When two different markers meet at the top of a hill, this is defined as the watershed line or point that will separate the two markers. This is often, and best illustrated with a 1D example – see Figure 4.6. The explanation in 2D or 3D is no different in substance, only very much more difficult to illustrate. Note that the entire 1D image is assigned either to one of the three markers, or as a watershed point. This will also be the case in 2D and 3D: the expanded markers and watershed lines/surfaces (which are closed lines/surfaces) will fill the entire volume by the end of the application of the algorithm. From now on, watershed lines/surfaces will simply be referred to as watershed lines, for readability.

The watershed lines are the output of the watershed algorithm that will be used to separate the grains from each other. The voxels defined as watershed lines will be *subtracted* from the solid phase in the binary image, giving a binary, separated image. This means that particles are separated by deleting a surface of pixels between them.

The deletion of contact voxels between grains is not strictly necessary (Some implementations of Meyer (1994) for example define the delimitation between markers as between two voxels (*i.e.*, interpixel), without having a line with a finite thickness, see Section 4.2.5), however it is very convenient for the next step, in which particles are labelled. Furthermore, the subtraction of the watershed lines from the solid phase also means that the “precious” volumetric information of the solid phase (which has been selected to represent the measured volume of the solid phase), is lost. The voxels subtracted from the solid phase are not simply discarded: they are saved as contact voxels, which are treated later in Section 4.1.4.

Figure 4.6 illustrates the progression of a watershed algorithm on a fictitious (and smooth) 1D example, showing how markers are grown with respect to the landscape (which is 1D in this case). The landscape that will be flooded in the case of this study is yet to be defined: and there is a fair amount of liberty in the choice of this *segmentation criterion* or *function*. The objective is to have some (3D) landscape with troughs around the markers and in the centres of the grains, and with particularly high peaks where grains touch, in order to have an accurate contact. The landscape’s relief away the contacts is not so important, since away from the contacts the watershed line is expected to run in the void phase. This means that it won’t be subtracted from the solid phase, and therefore has no effect.

A potential candidate for a suitable segmentation function is an inverted EDM (which has already been computed for the definition of the markers). An has peaks (well inside grains) inverted to troughs for the segmentation function. Figure 4.3 b) and e) show together that this choice of segmentation function is likely to work well: any two markers in e) that are connected through the solid phase will have a watershed line (in 2D) drawn to separate them that runs along the minimum of the uninverted EDM shown in b).

Other possible candidates are the inverted greyscale image itself, or the gradient of the greyscale image. These both use the rich greyscale information in the reconstructed 3D image,




Figure 4.6: 1D example of the operation of a watershed algorithm, starting from 3 markers. Some key moments in the application of the algorithm are highlighted (1 to 7). Inspired by the clear explanation in Bernard *et al.* (2011)

which is encouraging since the inverted euclidean distance map proposed above is derived entirely from the binary image, and ignores these shades of grey. However it must be recalled that the main feature desired in a segmentation function is that of high ridges at the contacts, since it is here that the separation lines must be defined. The inverted greyscale image has some of the desired properties: the grains themselves are troughs (although they are a little rough due to the greyscale character inside the grain), and the pore is all high. A problem appears however at the contacts: if one considers a contact which is not punctual, but rather a contact between surfaces (such as the very flat contact illustrated in Figure 4.4 a), then there might be a few pixels which are all grain, the pore being too far away to have a noticeable partial volume effect. This condition opens the possibility for a connection from one marker to the other, all of which is at the greyvalue of the grain. This means that there is a risk that between two touching grains, even with a very low immersion depth, the markers will fill the grains very rapidly, thus with the potential for error. Subsequent increases in immersion of the algorithm would be spent slowly and uselessly climbing from the inside of the grain to the pore phase through the partial volume voxels.

The gradient image is also not suitable as a segmentation function for similar reasons: the algorithm relies on the gradient of the reconstructed 3D image to generate peaks at the edges of grains. As above, this works very well at a solid/void interface, but at a contact, the gradient information is degraded (with respect to the clear gradient at the solid/void interfaces), and so just where it is most needed, this technique falls down.

One potential solution to finding an improved segmentation function, which would allow the direct use of the rich greyscale images, or gradient images, are segmentation methods using implicit surfaces and level set methods (*e.g.*, Osher and Fedkiw, 2003). These are being investigated in a joint project with Caltech (see Vlahinić *et al.*, 2012). Very briefly, the advantage of these segmentation methods is that the boundary of the markers (when considered in a mechanical analogy) can be assigned some parameters, so that as it grows (again based on a function like the *segmentation function* seen above) it cannot deform into any shape: for example a limit on the maximum curvature could be imposed. This helps the marker respect the real grain’s geometry to some extent. Furthermore, these segmentation methods may yield (due to the continuous nature of the level sets’ surface) sub-pixel placement of the delimitation between grains, which may greatly increase the quality of the information relating to the contacts between segmented grains.

The inverted EDM is used by Visilog[®]’s “binseparate” marker finding and watershed command (Bernard *et al.*, 2011), which is the separation algorithm selected towards the beginning of this work for its speed – particularly important for the large volumes of data collected.

Some errors in separation cannot be avoided – as shown in Andò *et al.* (2012a), and these are show in Figure 4.5, and are a result of the kind of error in marker placement discussed with Figure 4.4.

Section 4.2.5 contains some further comments about the separation process, due to the poor performance of the techniques described here when the watershed lines are considered as contacts, and their orientations are studied.

4.1.3 Labelling

Once particles are separated in each image, a very simple program can traverse the 3D volume, assigning all voxels that are in contact with each other (*i.e.*, all voxels within a watershed) a single, unique value. A program of this type typically does not make any rational choices in the values to assign: increasing values are assigned as the program traverses the volume (typically from top to bottom). The result of this is a 3D image in which all the voxels of each particle have a unique label (a positive integer) assigned to them. This in turn means that a given particle can easily be extracted from the labelled 3D volume by selecting only voxels corresponding to its value. Code extract 4.1 contains an example in PYTHON (with SCIPY and NUMPY) showing

Code 4.1: Specimen Python code for the labelling of a separated volume, followed by the extraction of the position of the highest-numbered particle

```

# The 3D array "image_sep" contains a binary image, which has been
# previously separated with a watershed algorithm
import numpy
from scipy.ndimage import label
from scipy.ndimage import find_objects

# The labelling procedure, as implemented here, returns a labelled
# image and the highest label assigned.
[ image_labelled, highest_label ] = label( image_sep )

# Let's get the coordinates of the bounding box around the
# particle whose label is "highest_label"...
coordinates = find_objects( image_labelled == highest_label )[0]

x_start = coordinates[0].start
x_stop = coordinates[0].stop
y_start = coordinates[1].start
y_stop = coordinates[1].stop
z_start = coordinates[2].start
z_stop = coordinates[2].stop

# Extract a subvolume which contains the grain whose label is "highest_label"
subvolume_labelled = image_labelled[ x_start:x_stop, \
                                     y_start:y_stop, \
                                     z_start:z_stop ]

# In this subvolume keep only voxels that are equal to "highest_label".
# numpy.where's syntax is:
# numpy.where( condition(array), if_true(array_element = X), if_false(array_element = Y) )
subvolume_labelled = numpy.where( subvolume_labelled == highest_label, highest_label, 0 )

```

the ease with which a labelled 3D image can be handled, and the coordinates of the bounding box around a labelled particle extracted.

Labelling allows, therefore, the definition of the extent of each particle in a given state. It is important to recall that this definition of the particles comes from a binarisation and a separation of the original greyscale 3D image coming from reconstruction. Consequently, if the grey information pertaining to the grain is desired, it can be obtained by “looking up” the greyscale 3D image, using the labelled image as a mask (*i.e.*, keeping only the information in the greyscale 3D image that corresponds to voxels of interest in the labelled volume). There is however a problem with this: the combination of partial volume effect (outlined in 4.1.1) and threshold means that only the highest information on the edges of the particles will have been selected, the lower partial volume voxels are not selected. In order obtain these voxels, the labelled particle desired can be *dilated* (*e.g.*, expanded by one voxel all over), increasing the size of the mask, and allowing the lower greyscale information to be captured. This is done with the risk of “invading” other grains. This is illustrated below in Figure 4.7.

Looking forwards to Chapter 5, using the process described above, there is no guarantee that labels will be assigned to the same physical grain between configurations. This is because a labelling program simply traverses the volume and labels particles in the order they are met. However, if a list of different measurements for labelled particles is available, a technique like ID-Track described in Chapter 5 can be used to follow grains from one configuration to the next. ID-track simply matches labels in one configuration with labels in a following one, attempting to recognise grains locally based on a single measurement (3D volume has typically been used). It is described in the following chapter in Section 5.3.




Figure 4.7: Flow chart showing how the grey information for a given grain is obtained. With no dilation, only the grey information on the inside of the grain is copied. With one cycle of dilation, the darker partial volume voxels are copied. With too much dilation (two cycles are sufficient in this case, grey information belonging to other grains is copied

4.1.4 Definition of contacts

Contacts between grains are the points at which forces are transmitted in a granular assembly. The characterisation of the existence, position and orientation of contacts between grains is therefore a key measurement to be able to make on a granular assembly. The presence of contacts between grains is relatively easy to establish: it is a natural output of the separation process (Section 4.1.2). The voxels that the algorithm establishes as being at the delimitation between two particles are deleted in order to separate the solid phase into individual particles. If the separation has been performed correctly, the voxels that have been deleted in order to separate two particles represent themselves the position of the contact in space. It is unfortunate perhaps that the contact between two particles is made of voxels and therefore has a volume. As briefly discussed in Section 4.1.2, this is not necessarily the case (intervoxel watershed or separation lines are conceivable). However in most watershed implementations, watershed lines have a unit thickness.

Contacts are defined in this work as the sum total of all the points at which two particles touch: so if two particles touch each other in two places, this is considered as *one* contact. In order to be able to individuate and measure individual contacts, as with the particles described above in Section 4.1.3, they need to be labelled. Two complications render the process of labelling contacts more subtle than the labelling of grains discussed above: first, voxels representing

the same contact need not necessarily be contiguous in space, and second, for a contact to be labelled, the two grains that touch at that contact need to be recorded. Because of this, a special contact labelling code was developed in the context of this doctoral work to overcome these two complications. It is outlined in flow chart form in Figure 4.8.




Figure 4.8: Flow chart describing the procedure for labelling contacts between grains

The code is not immediately straightforward because, for an unknown number of contacts, each contact voxel touching the same two grains, must be given the same label. Reasoning from the fact that we know that there approximately 50,000 grains in a given 3D image (for the sake of this example we will say there are exactly 50,000 grains), and that we expect the number of contacts per grain to be around 10 (*i.e.*, not 1 and not 100), an empty “main” list with 50,000 elements is created. The idea is to store each grain contact based on which grains it is contacting – a grain-based reference system. Pre-assigning this list from the grains’ point of view is beneficial, since the grains are already labelled, and so avoids the list growing unnecessarily in memory. The list *will* grow in memory, as contacts are labelled, but pre-assigning the list of grains is most *a-priori* information that can be pre-allocated. This algorithm, for every contact voxel, looks at the immediately neighbouring voxels in space, and records the different *grain* labels present (as shown in Figure 4.8 a). If there are exactly two different grain labels found, then the smaller label is called A and the bigger label is called B (Figure 4.8 b). The following procedure is then followed (as per part c of the figure):

The main list’s Ath element is interrogated. If main list’s Ath element contains any lists, the first element of each list is compared with B. If B is found this means that a *contact label* has

already been assigned to the contact between grain A and grain B. This contact label is read from the second element of the list in which B was found. This contact label is then assigned to the voxel in question.

If B is not found, this means that a new contact has to be defined. This is done by incrementing the “contact count” by one, and by adding a two-element list (containing B and the contact count) to the Ath element of the main list. A corresponding list [A, contact count] is also added to the Bth element of the main list – see Figure 4.8 The voxel in question is assigned the new value of contact count.

At the end of this procedure, the main list is processed into a 2D array, which contains sufficient columns to allow the copying of the element of the main lists which contains the most number of two-element lists (*i.e.*, the grain which is in contact with the largest number of other grains). This number needs to be determined before the processing into an array, and explains why a big array is not used from the beginning: the maximum number of contacts is not known *a priori*. Importantly, a further output is created: a list, for all the grains, of the other grains that it is in contact with. The array is saved as a grain-contacts file, where the rows are monotonically increasing grain labels (from 1 to 50001 in this case) and the columns specify “Contacting Grain Label N” and “Label Given to this Contact” progressively from N=1 to N=largest single number of contacts per grain.

The overall result is a labelled contact volume, and a table of grain to contact labels. The voxels belonging to a contact can be extracted with a bit of code, just like the one shown for grains in the code extract 4.1. The fact that voxels are non-contiguous is not necessarily a problem.

4.2 Measurement of grain and contact properties

So far in this chapter, the image analysis required for the definition of individual particles in images coming from greyscale 3D images has been detailed. As seen in Section 4.1.3, and in particular in code extract 4.1, from a computational point of view it is relatively easy to extract particles individually from a labelled 3D image. This being the case, a lot of different properties can be calculated for each particle, using the contiguous set of voxels that represent the particle. For example, the volume of a particle can easily be calculated by simply counting the number of voxels that make it up. This in turn can be converted to a volume in mm^3 knowing the voxel size in the image. From this definition of particles, the following voxel-based measurements are forthcoming:

- Centre of Mass (often simply the mean of the positions of the voxels in each direction)
- Orientation (typically calculated from the minimum and maximum eigenvectors of the moment of inertia tensor of the particle)
- Surface Area (counting voxel faces that touch the void phase)
- 3D Volume (number of voxels making up the particle)
- “Feret” Measurements (minimum and maximum widths, measured over a certain number of different directions)

These different measurements made on these 3D binary objects represent some fundamental micro-scale measurements. The measurement of the position of each particle (using its centre of mass) is of particular interest, since chapter 5 will attempt to follow particles from one configuration (or image) to another. If particles can be recognised between configurations, their displacement can be measured, and therefore some characterisation of this measurement is important.

The orientation of particles is also of great interest; however, results obtained using these measurements have proved to be inadequate, so these measurements are not characterised fully. However an outline of the reasons of their inadequacy is given in Section 4.2.2.

Other measurements are not used to such a profound extent as the measurement of positions, and so are briefly described and are taken as given from the measurement toolbox in Visilog[®] and/or PYTHON– NUMPY and SCIPY.

4.2.1 Measurement of particle positions

In this work, labelled particles' position is measured by calculating the centre of mass of the voxels that make up the particle, giving its position in the coordinate system of the image. The simplest approach (and the approach used in this work) is to affect the same mass to every voxel belonging to this particle and calculate the centre of mass as the spatial average of the totality of voxels. In all that follows, the centre of mass of a pixel or voxel is considered to be in the middle of the pixel or voxel: so pixel located at $(0, 0)$'s centre of mass will be $(0.5, 0.5)$, in the style of Glassner (1993).

Assigning error bounds to this measurement of particle position is not trivial, because in discretised (voxel) space the sensitivity of measurements is dependent on the shape, size and orientation of the particle in question. The fundamental source of error is the combination of the discretisation of space into a 3D matrix of points, and the averaging of the measured properties of space inside each point. The latter is sometimes known as the partial volume effect (outlined in in 4.1.1) whose name recalls that if the space imaged by a single voxel is shared between a grain and air (partially occupied by a grain), the resulting measured property (reconstructed x-ray attenuation coefficient in this case, or “CT-value”) will be an average of the two materials present in voxel's space. Because of the partial volume effect, the 3D volume resulting from an x-ray tomography on sand grains measures the expected x-ray attenuation coefficient for sand at the voxels well inside a grain, and measures the expected value for the pore fluid (air in the case of these experiments) well inside the pores. There is a zone, a few pixels long, of “partial volume voxels”, which have intermediate values. When a correct threshold value is applied to such an image in order to select the grains, some of the “more dense” partial volume voxels will be selected along with the interior of the grain. The effects of this binarisation of the image into grain and void are the necessary creation of jagged edges. This is best studied with a simple 1D example.

In one dimension, assuming that the partial volume effect works linearly (*i.e.*, that a voxel's x-ray attenuation value is directly proportional to ratio of grain and air in the space imaged by the voxel) some simple examples of the displacement of different 1D bars reveal some fact about measurements where image binarisation is concerned. The Figure 4.9 illustrates the concept of binarisation, and underlines the “harshness” of the threshold which is applied to the image. As can be seen, unless a pixel is darker than 50% grey (*i.e.* the 1D pixel is more than half occupied by the shape) it will not be selected by the threshold in the binary image. Because of this, if as is the case below, shapes are aligned left, there is no difference between the binarised image of a line 3.501 pixels long (pixel 4 is just over the threshold) and that of a line of 4.499 pixels long (pixel 4 well over the threshold and pixel 5 just under the threshold) – they are both 4 pixels long.

A very simplified version of the standard approach to characterisation the metrological aspects of a measurement (as described in Greve (1967)) would require the measurement of the sensitivity of this measurement, and following this, a characterisation of the precision and accuracy of the measurement.

The conclusions that can be drawn from this 1D example above is that in an equally weighted binary image, the centre of mass of each line of pixels can be defined with maximum error of 0.5 px per edge.

It is worth noting in this example that it is assumed that the grey image can have any value

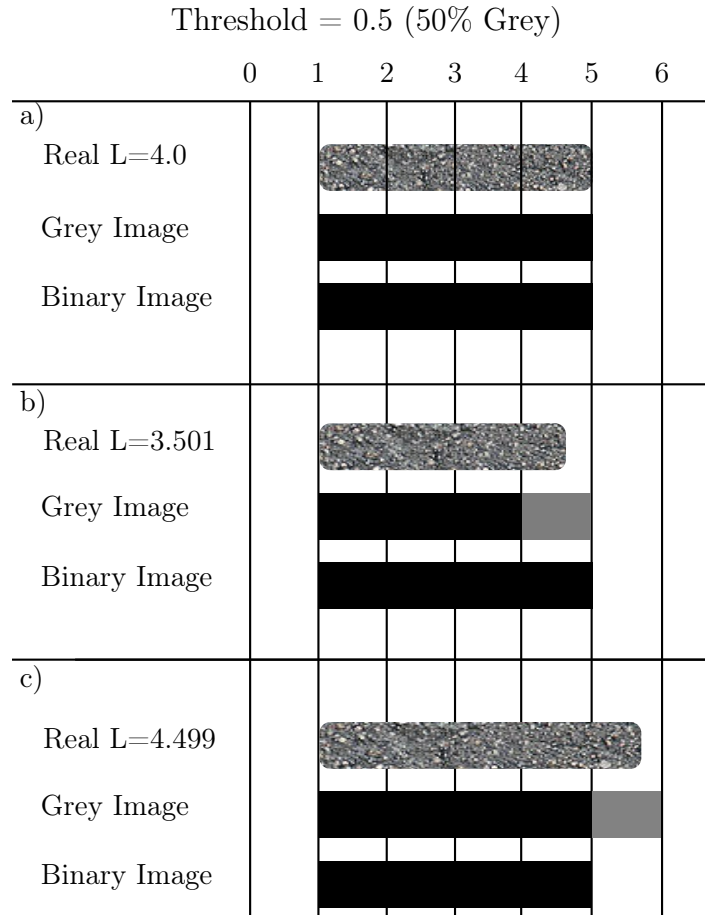


Figure 4.9: Simple 1D example showing 1D lines of different lengths (in pixels) being imaged by a greyscale detector, and then thresholded. All the lines of different lengths result in a binarised line 4 pixels long, after thresholding

of grey (*e.g.*, 4.499) – and does not have a minimum sensitivity. In reality all digital images also have a sensitivity associated with the measurement of each pixel. In this experimental work a 3D x-ray attenuation coefficient is reconstructed from more than 1000 radiograms, inside which each pixel's value has 14 bytes of storage ($2^{14} = 16384$ different possible values), so the resulting image has an extremely low sensitivity threshold with respect to the grey value of each pixel or voxel.

In this work, the centre of mass is calculated as the spatial average of all voxels in each direction (since the voxels have equal weighting). The centre of mass of a 1D object, whose image or mass distribution is $I(x)$, beginning pixel a and ending pixel b takes the following form:

Defining the length of the line:

$$l = b - a + 1$$

$$COM = 1/l \times \sum_{x=a}^b (I(x) \times (x + 0.5))$$

Note this also works for greyscale images. If, as stated above we elect to assign the same mass to each binary grain pixel. The mass of the pixels becomes redundant in the calculation, giving:

$$COM = 1/l \times \sum_{x=a}^b (x + 0.5)$$

Making a sum over the length we now get:

$$COM = a + 1/l \times \sum_{x=a}^b (x + 0.5)$$

$$COM = a + 0.5 + 1/l \times \sum_{x=a}^b (x)$$

Since the sum of an arithmetic progression of l numbers starting with a , ending with b , and with difference between successive numbers equal to 1 is given by:

$$Sum = l/2(a + b)$$

Substituting:

$$COM = a + (1/l) \times l/2((a + 0.5) + (b + 0.5))$$

Simplifying:

$$COM = a + (1/2)(a + b + 1)$$

Since a and b are necessarily integers, this development shows that the sensitivity of the measurement of the centre of mass is also limited to 0.5 pixels for binary images, and also importantly, that the centre of mass calculated in this way depends only on position of the pixels on the border (a and b).

This result shows that when calculating the centre of mass of a binary line, *only the position of the front and back pixel (a and b) are important* (or more generally: only the perimeter of the shape is important), since the mass of pixels inside the shape plays no role in moving or weighting the centre of mass. It is important to note that this is not the case for greyscale images, where the pixels on the perimeter may be a certain shade of grey, and therefore have a lower mass – and consequently be outweighed by the mass inside the shape. A partially-filled grey pixel's own centre of mass is in the middle of the pixel, so this skews the centre of mass, unless there is a corresponding skew on the other side: it is for this reason that the centre of mass of a line whose length is an integer number of pixels is always correct.

To sum up: in one dimension, the maximum sensitivity and therefore maximum error in the calculation of the centre of mass of a binary image, with respect to the real position of the object being imaged is of 0.5 px, regardless of the length of the line. With a 1D greyscale image, the sensitivity and error decrease as the line's length increases (at integer lengths the error is 0). This is illustrated in Figure 4.10.




Figure 4.10: Graph of the error in the calculation of Centre of Mass, as a line grows from 0 to 10 pixels in length (One end of the line is stationary at pixel 0). Note that there is no centre of mass for the binary image until the line reaches 0.5 px in length, where the first pixel is bright enough to be selected by the threshold. Note: The error here only goes as high as 0.25 pixels, since the bar grows only from one direction

Extrapolating from this 1D example, to a second dimension, we can identify a worst-case shape and alignment and a best-case one. If one considers a set of 1D rows of pixels as above, all the same length and all aligned along the axis opposite to their length, we have a square. If we displace the square in the direction of the bars, the reasoning for the single bar works for a group – there will be a value (or two values if the bar is not an integer number of pixels long) of displacement which will make the all 1D bars’ end pixels change together. This means that the worst-case sensitivity and maximum error for the measurement of the centre of mass will be the same as the 1D bar. The best case shape would have to have a range of different length 1D bars, spaced in such a way that in the continuous range of sub-pixel displacements from 0 to 1 pixels in the direction of the bars, the end pixels of different bars flip at evenly-paced intervals. Each pixel that flips would gently shift the centre of mass. A 2D shape made of 6 bars each having one end spaced at 1/6th of a pixel from the next one is the best shape that width for detecting movement in the direction of the bars. In this case it is clear that the best case (or minimum bound) for the sensitivity in this case is:

$$\text{Best Case Sensitivity for 2D Shape} = \frac{\text{the sensitivity for one bar (0.5 px)}}{\text{the number of bars that make up the shape}}$$

This goes to show that the shape being imaged is important, since a square aligned with the pixels represents the worst-case sensitivity, and a quadrilateral shape with a taper represents the best-case sensitivity to displacement, when binarised.

Analytical results

In the literature that deals with this kind of problem (the precision of measurements made on digital images), a lot of attention has been paid to finding analytical results for simple, 2D geometrical shapes known a-priori. Examples include pixelised lines (*e.g.*, O’Gorman, 1991), points (Havelock, 1989), circles (Landau, 1987; Nakamura and Aizawa, 1984), this work has often been carried out with an application for computer vision for automation. These references

provide some analytical results for measurements made on images of these shapes. Because of the predictable nature of the pixel patterns that surround straight lines, and the known symmetry and dimensions of circles there is a large quantity of a-priori information that can be integrated into a measurement.

A potentially useful idea proposed by Havelock (Havelock, 1989, 1991) is that of locales. Locales are a feature of a shape – very simply they describe the regions in pixel- or voxel-space in which a real shape can move while its binarised image stays the same. Continuing the development of this concept, and adding the calculation of the centre of mass, it may be possible to characterise the sensitivity, precision and accuracy with which a shape can be placed in 3D. This is highly desirable, but requires “perfect” geometrical knowledge of the shapes under study.

These analytical solutions, and the concept of locales are unfortunately not suitable, at least in their current form, for the characterisation of the measurements made on the particles treated in this work for three main reasons:

1. The natural grains studied in this work are not perfect, symmetrical geometrical shapes
2. Each grain is individual (there is no accurate and generalised description of grain shape which can be given to all the grains of Hostun sand for example)
3. A variety of different grain shapes is used in this work. With little or no quantitative information about the shape of the grains under study, it is very difficult to put in place an analytical development of the behaviour of these 3D grain shapes when imaged, since plane sections in a shape can have an adverse effect on the accuracy of the measured centre of mass – especially if these sections happen to be aligned with the coordinate system of the image.

One possible way to approach some analytical solutions is to geometrically describe a number of grains, in 3D for each type of scan. This should ideally be done at a much higher spatial resolution than that of the measurement, in order to be able to describe the shape of each grain with high precision. One such approach is proposed in Vlahinić *et al.* (2012) and hopes to use level-sets to segment grains along smooth boundaries, allowing them to be described by smooth Non-uniform Rational B-Splines (NURBs). An analytical study of such a shape may be hard, however a grain described in such a manner could be numerically displaced and binarised, allowing locales to be calculated for it, which would give a good indication of the errors possible measuring the displacement of such a grain. Since these grains are in no way guaranteed to be symmetric, it is important to consider displacement in all directions or orientations.

Numerical simulation

In order to simulate the sensitivity, error and precision of the measurement of subpixel displacements of different geometrical shapes, a very simple program was written in which simple geometrical shapes can be defined (lines, circles and quadrilateral shapes in 2D, and their equivalents in 3D). These shapes are “imaged” in greyscale onto a matrix of pixels (by sub-integrating the shape inside each pixel – the sensitivity of the greyscale depends on the number of subpixel integration points chosen). This image is then binarised, with a threshold half way between the value of the shape and the value outside. The centres of mass of the binarised image, the greyscale image and the originally imposed shape are then compared, by looking at the deviation of the greyscale centre of mass (which is calculated by weighting each pixel by its greyvalue) and the binary centre of mass from the real centre of mass which has been imaged. This is done for a variety of different sub-pixel displacements.

A 1D simulation (as considered earlier in this section) is performed in order to verify the results found – in this case a 1D bar with an integer length pixels is displaced in the direction of its length. Results are shown in Figure 4.11.

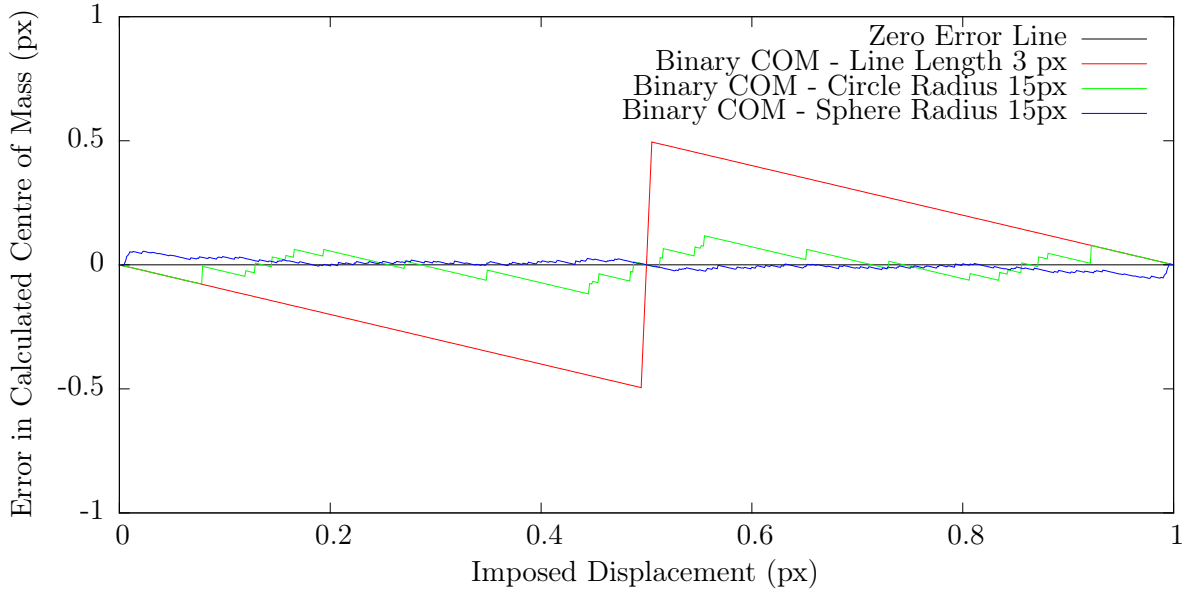


Figure 4.11: The error in the measurement of centre of mass of different shapes, when they are binarised. Each shape is displaced from 0 to 1 pixels with small subpixel increments in the x direction

As developed above, this result confirms that the binarised image’s sensitivity to this line’s displacement is 0.5 pixels. As the line displaces, the error in the binary centre of mass increases, since the centre of mass of the shape being imaged is actually moving, whereas the binary image has no change until it has displaced 0.5 px. When the line has displaced 0.5 px, the leading pixel, which has got progressively brighter due to increasing occupation of the front of the line into it gets to a value of 0.5 px. Similarly the trailing pixel, which has got progressively darker gets just under 0.5 px. At this point, the binarised image “snaps” forwards 1 px, by losing the back pixel and by gaining the front pixel. Relative to the shape being imaged, the distance between the error in the measurement of the centre of mass has moved from -0.5 px to 0.5 px. Just before (and just after) the binarised image changes, the binarised shape’s centre of mass is at the maximum error with respect to the real image, and this is at 0.5 pixels (for any length of bar).

Using the same logic as above, if a block of several of these lines was displaced in the direction of the lines, the centre of mass, in the direction of the movement would behave in the same way.

The displacement of a circle is also simulated in Figure 4.11. The circle is chosen to have a radius of 15 pixels, approximately like the cross-section of grains in a slice. The error in the measurement of the centre of mass with respect to the displacement imposed has a maximum error which is much smaller (maximum error in the simulation shown, displacing the shape 1 pixel in 1000 steps was 0.1173 px). This error is far reduced because of the geometry of the problem. The fact that the circle is not aligned to the coordinate system of the image, means that there are a lot of partial volume (or more correctly partial area) pixels on the border. The partial area pixels are in the order of 118 pixels in this case, compared to the two in the case of the line. When this image is binarised, the fact of having many partial area pixels means that for a small displacement there is an increased probability that a pair (due to the symmetry of the shape) of pixels will change phase. A pair of pixels that flip will have a small effect on the centre of mass, which will move as a result. Using the 1D analysis above (and taking into account only half a circle because of symmetry) the best-case sensitivity to displacement (which would occur if the partial area pixels represented a continuous variation of grey values) would be $0.5/15$ pixels which is 0.0333 pixels – 3.5 times less than the maximum error measured above.

Continuing to a 3D sphere of radius 15 pixels, also present in Figure 4.11, one can see that

the error in the calculation of the centre of mass is even smaller. In this case the sphere has in the order of 4256 partial volume voxels. This gives a maximum error (again when displacing the shape 1 pixel in 1000 steps) of 0.0547 pixels. In this case, the best-case error, the number of unique 1D bars is the number of 1D bars on the widest diameter divided by four (for symmetry). One quarter of the area of a circle radius 15 pixels is 176.71, the simulation gives 179 pixels above the threshold. The best-case error in this case is $0.5/179 = 0.0028$, i.e. 20 times better than the simulation. As can be seen in Figure 4.11, the sphere’s maximum error lies close to the extremities.

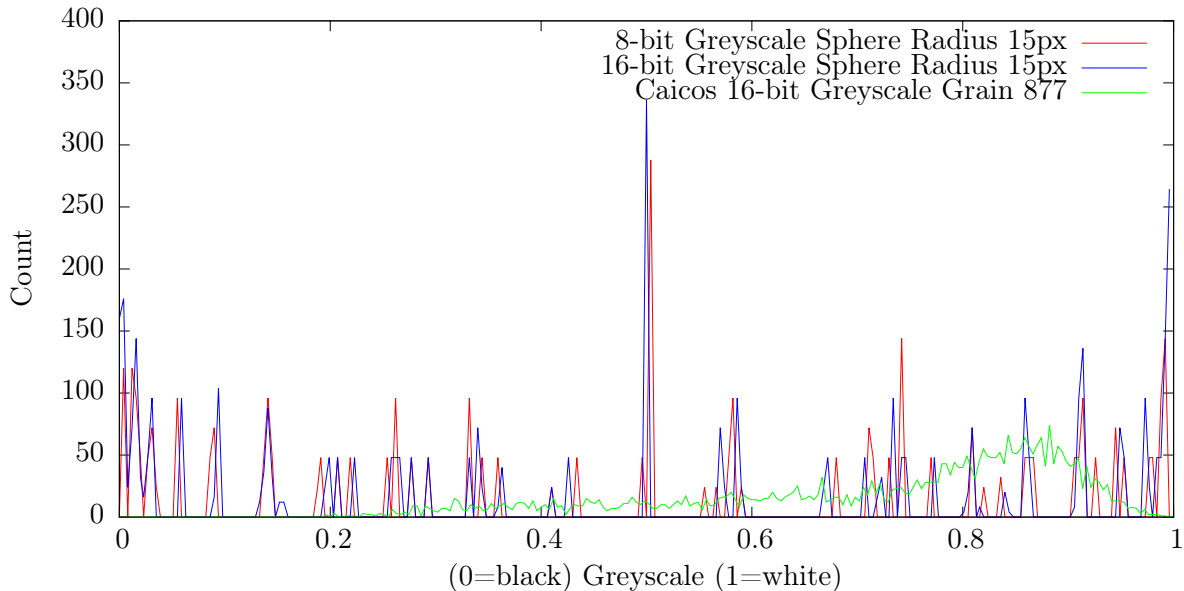


Figure 4.12: Greyscale histogram of the partial volume voxels of the sphere coming from the first step of the simulations described. Shown for 2^8 or 2^{16} subintegration steps inside partial volume voxels. Values of “count” for values 0 and 1 are hidden to show the partial volume voxels

Figure 4.12, spheres, show that the spread of different grey values, in the first step of the displacement simulation is not very equal. This may explain the considerable difference between the best-case error and the actual maximum error recorded, as well as some of the trends visible in the sphere’s graph in Figure 4.11. The existence of these very well-defined peaks in the image are due to the perfect geometry of the system.

The grey will be spread much more homogeneously in a real greyscale image. This is difficult to show in the same format, since there is not a single greyvalue which represents the grain or the pore. Consequently, the entire histogram of one grain has been plotted on 4.12. It is clearly visible that for the real grain, the histogram is spread much more equally. This means that the threshold that selects partial volume voxels will do so much more sensitively, and in less of a blocky fashion than the threshold in the case of the sphere.

Effect of watershed process

The effect of the watershed’s removal of voxels at the contacts between grains has not been considered in the developments so far: particles have been considered so far alone. The removal of voxels at contacts (with watershed algorithms that do not giver interpixel watershed lines, like the one used in the main body of this work) necessarily deletes some of the solid phase. However, the effect is very small: for the test COEA01, the mean number of voxels making up a contact is 51.2, which means about 25 voxels lost per grain (for each contact). Compared to the 6761 voxels that on average make up a Caicos Ooid grain, this is a very small quantity indeed (0.37%).

Furthermore the fact that grains will have multiple contacts arranged typically at opposite ends (for mechanical stability) will tend to reduce the effect that the removal of these voxels will have. For example, a grain's centre of mass will not change if it has two contacts, one top and one bottom, since the effect on the position of the centre of mass of one voxel is cancelled out by the other one.

Proposed methods of improvements

To increase the quality of the measurement two principal strategies are possible – either working to improve the binary image, or using the greyscale images from which the binary images were originally derived. The principal advantage of improving the binary image is that one can work on the original definition of the particle in the image without having the added complication of using the separately-saved greyscale information. One approach may be to improve the binary image by rendering it less jagged, this could be done either by making a subpixel refinement of the boundary pixels using the Marching Cubes Algorithm in 3D (Lorensen and Cline (1987)), which is commonly used to smooth voxelised objects for 3D rendering, or by 3D curve-fitting the voxels of the particle, as discussed above. Both these approaches will smooth the unrealistically jagged shape of the particle, which will help to render it more physically correct, however for very angular grains there is the risk that correctly images sharp edges are lost. The other possible approach, which would be much simpler to implement is a weight to the binary pixels. Since the most uncertain voxels are those on the boundary, it is possible to automatically reduce the weight assigned to the partial-volume voxels on the outside of the particle. A region two or three pixels deep should be affected by this reduced weight, reducing the weight progressively as one gets closer to the surface of the shape.

Improving the precision of the measurement with greyscale image is also possible – for each particle, it is necessary to look up the positions of the voxels that make up the particle (defined in a binary or labelled volume – this is illustrated in Figure 4.7) in the original greyscale reconstructed volume. This looked-up 3D shape includes a certain quantity of more grain-like partial volume voxels and so has a softer outer edge than the binary image. The more pore-like partial volume voxels that were not included by the threshold will however not be collected in this manner, so in order to get those as well it would be necessary to dilate the original particle by one or two voxels in order to get all the greyscale voxels that represent the grain. The risk in this case is that with dilation one may invade another grain, which would add incorrect voxels to the particle. To avoid this the dilation should be stopped at contacts. The centre of mass of the particle can be computed using the reconstructed x-ray attenuation coefficient of each voxel as its weight. From numerical simulations, computing the centre of mass this way is shown to be a lot more precise than calculating it from the binarised versions of the same data.

Conclusions

The conclusions of this investigation into the characterisation of centres of mass of binary objects are perhaps not very strong, due to the different types of natural materials studied. Since many different shapes of grains are dealt with, no *a priori* geometrical information is known on the grains, making a geometric analysis impossible.

When the concepts above were presented at a workshop of the metrology-specific subgroup of the (French) National Center for Scientific Research's "Full-field Methods and Identification in Materials" group (GDR-2519), the general conclusion from the discussion that followed was that, with such diverse natural grains, and with grains measuring about 15 voxels in radius, in the images there are probably a large enough number of partial volume voxels in the image, that small displacements can be detected.

The numerical simulation on the sphere radius 15 pixels performed in 4.2.1 gave about 4200 partial volume voxels. Remembering furthermore that in the images acquired, the partial volume

effect is spread over 3 pixels, a little wider than the very tight partial volume effect in the numerical simulation (only two pixels thick).

This large number of partial volume voxels are considered beneficial to the error and sensitivity in the measurement of positions by calculation of a centre of mass. For a very small sub-pixel displacement, in a large amount of partial volume voxels there is reasonable hope that a few voxels will change phase and thus add a small contribution to the centre of mass.




Figure 4.13: Radiograph of a single Hostun grain (glued onto the head of a toothpick). Taken with a nano x-ray tomography demonstrator courtesy of RX-Solutions

For rectangular shapes with the bad luck of being aligned with the coordinate system, the maximum error remains 0.5 voxels. However, since the grains being studied in this work are all far from being parallelograms, and even angular Hostun grains (as shown in Figure 4.13), are far from having many parallel lines, and are deemed not really to suffer from the alignment problems possible with parallelograms.

In conclusion, it is likely that given the slightly wider partial volume voxels, meaning a better distribution of greylevels going into the binarisation process, the maximum error *in each direction* of the centre of mass measured on these binary grains is better, but possibly in order of the sphere's 0.05 pixels.

At a spatial resolution of 15 μm as is typical in these images, this corresponds to a maximum error of 0.75 μm , which for grains with a D_{50} of more than 300 μm gives a good resolution.

Evidence in the results of grain tracking on very small increments will confirm this in chapter 5.

4.2.2 Measurement of particle orientation

The measurement of particle orientation is of obvious interest: allowing the characterisation of the initial state of a granular assembly, using micro-mechanical tools such as the fabric tensor. Furthermore, with the grain tracking detailed in chapter 5 the measurement of orientations between configurations would allow the measurement of rotations of grains.

A common technique for orienting particles (available, for example, in Visilog[®]) relies on the moment of inertia. The moment of inertia tensor (which allows the computation of the moment of inertia about an arbitrary axis) is typically calculated, weighting voxels equally, on all the voxels that make up a particle. The resulting symmetric tensor's eigenvectors define the

principal directions of the particle. The principal axis with the highest moment of inertia will point in the longest direction of the particle, from the centre of mass, whereas the principal axis with the smallest moment of inertia will point in the shortest direction.

For non-spherical grains (where the two measurements are confused) one might expect this to be sufficient, however two important problems arise from this definition of orientation – especially when this orientation is used computed for the same particle in two different configurations in order to measure its rotation (the procedure for measuring the rotations is illustrated in Section 5.3.5 of the next chapter).

The first problem stems from the fact that particles are defined on a binary image. When a grain is imaged and thresholded in two different configurations, it is possible (and even desirable – if the grain has moved) that between configurations some of the grain’s voxels change – this may reflect a small displacement for example. However, if one of the principal axes is not defined very uniquely for a given grain, the change of a few voxels may be enough to cause a different principal axis to be found (which would cause a non-physical rotation). This can affect both principal axes, and if one or both principal axes jumps in this way, this causes non-physical rotations to be measured. This source of error is difficult to quantify since it depends profoundly on the shape of grains being imaged. Even if shape is accounted for, it would be impossible to associate an error bound to this measurement, since a jump of any angle is in theory possible (and becomes increasingly possible as particles approach a sphere).

The second pitfall with these measurements is that the principal axes are just axes and not vectors. If the vectors that represent these axes are negated, the negated axis remains a valid principal axis for the given moment of inertia tensor. This means that the occasional 180° flip in either or both vectors can occur. When using these orientations to make measurements of rotations between steps, this means that the biggest rotation that can be measured this way is 90°. Furthermore, grains that are not moving may experience flips, causing them to register a large rotation. If only one of the two principal vectors used to construct the base flips, this will cause a 180° rotation to be measured. Because of this, rotations are “folded” around 90° so that a 187° rotation is corrected to a 3° one.

In any case, keeping in mind these potential sources of error, these measurements can be used to characterise orientations of grains to some extent.

4.2.3 Measurement of particle volume and surface area

The measurement of particle volume and surface area is straightforward on a binary set of voxels. Volume is calculated by simply calculating the number of voxels that make up the particle. The threshold applied in Section 4.1.1 to separate the solid phase from the void is done in such a way to recover the grain volume measured after test. However the solid phase is (in the current implementation) split into separate grains (see Section 4.1.2) by removing a 1-voxel thick surface between contact grains, which causes the correct physical volume of the solid phase to be lost. In test HNEA01, in step 01, the solid phase is composed of 463,738,597 voxels. The number of voxels removed in the watershed process are “only” 14,668,050 voxels (3.16%). This means that a small negative bias will be introduced in the volume of the grains measured. As stated in Section 4.1.2, interpixel watershed lines are a potential solution to this.

The surface area is also derived from a simple calculation: the number of faces of each voxel which is in contact with the void phase. Due to the blocky nature of the voxelised surface of the binarised particle, this measurement will almost certainly be a large overestimate of the real surface area of the grain. In order to improve this measurement, the Marching Cubes Algorithm in 3D (Lorenson and Cline (1987)) could also be used in this context, in order to give the binary particle a smoother outer surface, a better approximation of the correct surface area, at this resolution.

4.2.4 Measurement of particle width and length

Particle width and length (*i.e.*, maximum and minimum dimensions) are computed using Feret Diameters, also known as caliper lengths. In 3D this involves measuring, with two parallel planes on either side of the particle the distance that separates the planes when they are both touching the outside of the particle. This is done for a certain number of different orientations of the planes. The largest and smallest diameters are measured in this way.

The accuracy of this measurement depends on the number of different plane orientations tested.

4.2.5 Characterisation of contacts

The labelled contact voxels resulting from the contact labelling procedure described in Section 4.1.4 (which handles the fact that voxels belonging to the same contact, *i.e.*, between the same two grains, may not all be contiguous), can also be measured and characterised with the tools detailed in this section.

As stated in Section 4.2.2, characterising the micro-mechanical state of a granular medium is of clear interest in this work, and the orientation of contacts is an important component in the calculation of a fabric tensor for granular materials.

Since contacts are assumed to be relatively planar objects, the principal axis of the moment of inertia representing the smaller moment of inertia (*i.e.*, the short direction) should describe the contact's direction satisfactorily. Due to the problem of "jumps" of the principal axis vectors ("flips" are not a problem in this case, since only one vector is of interest), a more simple orientation measurement has been put in place, which takes attempts to fit a plane to the cloud of points (*i.e.*, voxel positions) that define the grain. A plane is fitted to this cloud of points by Singular Value Decomposition, and the resulting normal vector is used as the orientation of the contact.

Projection of 3D orientations for plotting

The orientation of contacts, as explained above, is an important information that is useful to characterise at the grain scale. Since a contact is considered as an essentially planar object, its orientation is fully described by the normal vector that defines the plane. The visualisation of these 3-component vectors in a 2D plot is not obvious, but a coherent way to do this already exists in geosciences: stereoplots. This type visualisation uses only unit vectors in the direction of each vector that needs to be plotted. These unit vectors are then all made to start at the origin. In this way they define a unit sphere, centred on the origin. This sphere is turned into a hemisphere by negating the vectors that are on the wrong side. The negative of a vector which defines a plane defines the same plane.

Since the experiments analysed in this work are performed on cylindrical specimens, which are therefore axially symmetric, the negative- z part of the sphere is discarded, and vectors pointing inside it negated.

This hemisphere is then observed from above so that describes a circle at its widest. This means that a contact whose orientation is vertical (its vector is positive z) points vertically, and so represents a point in the middle of the circle. A contact whose orientation is described by positive- x is found as a point on the perimeter of the circle, on the right.

The points on the hemisphere are then projected onto the plane with which the sphere was cut into a hemisphere. Many different types of projections exist, a simple one being a "straight down" projection, where the z -component of the points on the hemisphere is simply set to 0. This has the disadvantage of having a much greater density of points at the edges of the circle. To avoid this, the "Lambert azimuthal equal-area projection" is used. This preserves area (thus distorting angular relationships), meaning that for a set of points equally spaced on a sphere, the density of points in the projection is equal all over. It is for this reason that (for example

in Figure 4.14) the radial intervals are equally spaced: in a straight down projection high angle data would all be bunched up close to the perimeter of the circle.

This means that this projection allows any deviation from a uniformly spaced arrangement of vectors to easily be noticed.

Furthermore this kind of projection is particularly useful for “binning” the data, which can sometimes reveal trends that thousands of points plotted cannot. An equal-area projection allows binning of data even with bins of unequal areas (for example in bins of equal angles): since area is preserved in the projection, the density of projected points remains constant.

Orientations of contacts

Contact orientation is a key ingredient for micro-mechanical measurements of fabric. In this work, where irregular grains are used, contact orientations are defined as local normals to the grains in contact, and not the branch vector (which is the orientation of the vector which links the centres of mass of the grains in contact), which principally makes sense when dealing with spheres.

Figure 4.14 a) shows the orientation of a subset of the contacts in test COEA01, configuration 01, where Visilog[®] is used to separate the grains (and therefore to create the contacts). As stated just above in Section 4.2.5, the projections shown in Figure 4.14 have been done in such a way that axial symmetry is expected (assuming the specimen is prepared without introducing biases) in the arrangement of contacts. Axial symmetry in such a projection means that the projection should not vary with respect to the azimuth angle, *i.e.*, phenomena are expected to be concentric on this plot.

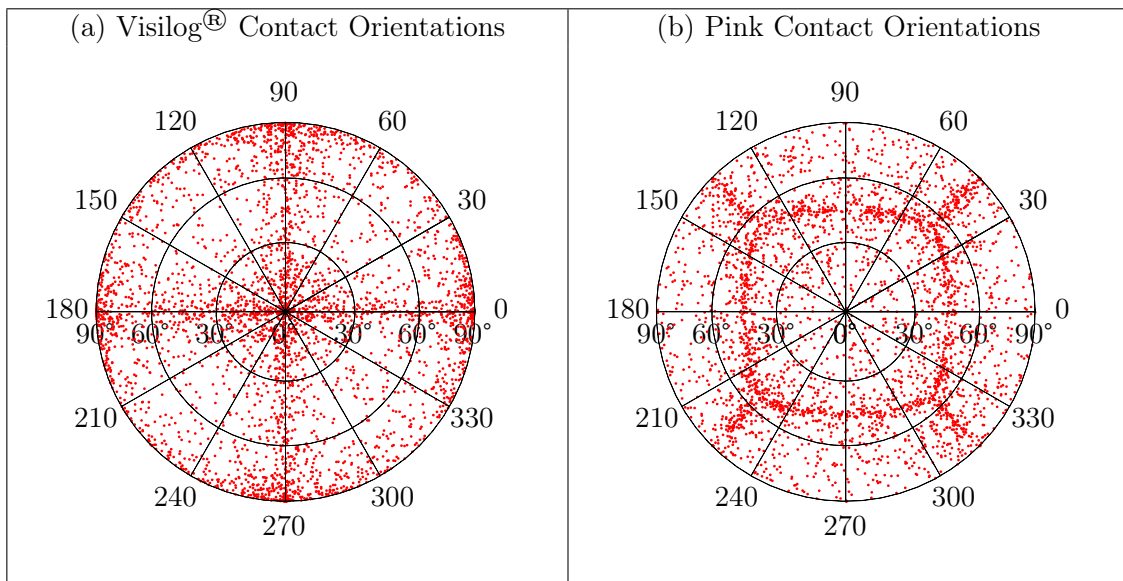


Figure 4.14: Contact Orientation for approximately four thousand contacts in COEA01-01. A Stereoplot is used to show 3D orientation in 2D with an area-preserving Lambert projection is used

Figure 4.14 a) shows orientations concentrated in a cross-like shape on this projection, which corresponds to a preferential contact direction of north-south and/or up-down and/or left-right, *i.e.*, aligned with the coordinate system of the image. This looks very suspiciously like an artefact, since it is aligned with the coordinate system of the image (the basis for all the calculations of the watershed algorithm) and is not constant with respect to the azimuth angle in the image.

In order to check this result, the contacts between the subvolume of grains published in Andò *et al.* (2012b), extracted from the same image (*viz.* COEA01-01), but segmented with the “pink.watershed” command from the pink image toolkit version 0.9.r646 (published in Marak

et al. (2011)). The orientations thus resulting are shown in 4.14 b), showing not only that the concentrations of points in a) are not only artefacts, but that the pink image watershed used also causes artefacts of a different kind.

This goes to show that the orientation of contacts as created by these two image processing libraries is both not reliable, as well as not inter-compatible. It is important to underline that despite the fact that these artefacts are a result of the watershed implementation, these artefacts are particularly aggravated in these images since the contact areas are small: Figure 4.15 is a histogram showing the dimensions of all the contacts between grains measured on images of Hostun sand, Ottawa sand and Caicos ooids. If contacts were much larger in the images analysed, the individual bias of each watershed implementation would have less freedom to manifest itself.

Figure 4.15 also shows the effects of the morphology of the grains, on the size of their contact area: the much more rounded Caicos ooids (red) have a clear peak in their distribution of contact sizes, showing that relatively round surfaces have a relatively well-defined contact area (at the resolution of observation – and when taking into account the partial volume effect). The fact that Caicos ooids also have some few irregularly shaped grains explains the tails to this distribution. The Hostun sand contact size distribution (blue) shows no peak for the dimensions shown – this is not surprising, for such an angular sand, many different types of contact topology are possible (in order of increasing size): from very small and extremely improbable point-to-point contacts, to the most likely (small) point-to-surface contacts to the much larger but improbable surface-to-surface contacts. The intermediate Ottawa sand does show a distinct peak for small contacts, corresponding to a point-to-surface contact, which is expected to be most likely – however given the more rounded nature of the grains, this has a larger characteristic area that the same configuration for Hostun grains. Similarly, when compared to Hostun sand, Ottawa sand grains have fewer large contacts, since fewer flat surfaces exist in these grains.




Figure 4.15: Histogram showing the distribution of contact sizes in imaged configurations COEA01-01, HNEA01-01, OUEA06-01 calculated with a bin width of 5 voxels

Towards better measurements of contact orientations

In order to investigate artefacts in grain orientations further, an intership is underway by Clara Jacquet at ESIEE, Paris under the supervision of Hugues Talbot and with regular input from the author.

The approach being developed to improve the measurement of contact orientations is to repeat the separation with a different approach, which offers a much richer information than the watershed line from classical watersheds. This approach coming from Grady (2006) and more recently Couprie *et al.* (2011) allows the calculation – given some initial markers – of the probability that a given pixel belongs to a given marker based on a random walk. Since this is quite computationally expensive (especially for large number of markers) the current implementation works on a single contact (and therefore two particles) at a time. Very briefly, the procedure basically consists in the calculation, for every voxel in the solid phase, of the random-walker probability of belonging either marker. An analogy used in Couprie *et al.* (2011) is that each pixel is connected to its immediate neighbours with a resistance equal to the gradient between the greylevels in the image, at that point. The preferential flow in this system can then be computed with a shortest path algorithm, for each pixel to each marker. Once this field is calculated, the voxels either side of the 50% probability value for both markers are identified; these are expected to be at the contact between the two particles. The position of the contact point for a pair of pixels (one with a marginally higher probability of belonging to one marker, and the other with a marginally higher probability of belonging to other other one) is then linearly interpolated based on the position of the pixels and their value of probability. This subpixel definition of all the contact points identified is then used as a point cloud, and fitted by Principal Component Analysis (see: Jolliffe, 2002). The fact that contact points are defined sub-pixel is expected to reduce the artefacts due to the pixels of the image.

An initial study was performed, consisting in the generation of many pairs of binary spheres (radius 10 pixels), contacting in a random orientation. Spheres are a particularly good test case, since the “branch vector” linking the centres of two touching spheres is in the same orientation as the normal to their surfaces at the point of contact. The markers used in this case are simply the centre points of the spheres. Figure 4.16 a) shows the imposed random orientations of the spheres calculated by their branch vector.

Figure 4.16 b) shows the orientations obtained if the spheres are separated using a classical watershed, with visibly the same artefacts as – recovering the same artefacts as 4.14 b). The situation is somewhat improved, in Figure 4.16 c), if a Meyer watershed is used. This type of watershed defines an inter-pixel separation between the two particles, however the orientations still retain a strong bias. However if the random-walker based technique is used, a much more uniform distribution of orientations resembling the distribution imposed is obtained.

The technique, having been qualitatively validated with this study, is then applied to the data coming from experiments. Since (for performance reasons) the random-walker based technique is implemented on a per-contact basis, grains need to be identified; to this end, the initial labelled volume and labelled contacts volume is used for this. The markers passed to the algorithm are the same markers used for the original segmentation. A first verification can be obtained on data acquired on glass ballotini (which are close to perfect spheres, and therefore branch vectors can be calculated – in this special case the branch vectors are considered as the most reliable measurement of contact orientation). Figure 4.17 a) shows the distribution of orientations of the measured branch vectors for a small subvolume of around a thousand glass ballotini. Figure 4.17 b) shows the orientations of the contacts obtained if the watershed lines from Visilog[®] are considered as contacts, and Figure 4.17 c) shows the considerable improvement obtained in the distribution of orientations with the random-walker based technique.

Given this successful verification, this technique has also been tested on data from non-spherical grains, where the orientations of the branch vectors are not necessarily the same as the normals to the grains surface at the contact points. Figure 4.18 a) and b) shown the distributions of the orientations of contacts (measured with the random-walker based technique), for a small subvolume of grains at the beginning and at the end of test COEA01. The subvolume is labelled “GL” and contains around a thousand grains studied is found in the shear band by the end of the test. An evolution in the orientation of contacts can finally be seen between the initial and

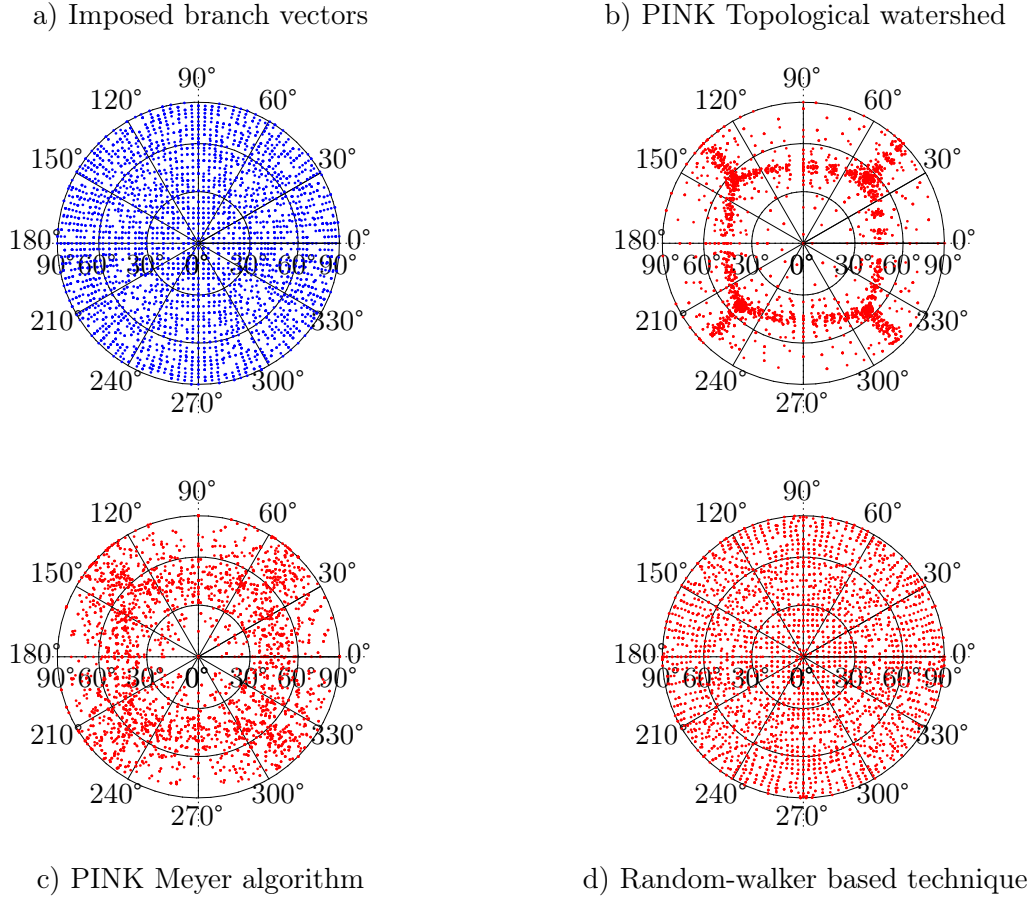


Figure 4.16: Equal area projection of contact orientations from simulated spheres

final states. In order to make this distribution easier to study, the projected orientations are divided into equal-area bins on the plane onto which they have been projected (Figure 4.18 c). Using cylindrical coordinates (r and θ), the circle in the plane is divided into equally spaced bins according to r . Radial bins are labelled $n = 0, 1, 2, \dots$ from the centre of the circle, and the number of angular bins in θ at a given bin distance is: $4(2n + 1)$. The data shown in figure 4.18 a) and b) is binned using 5 radial bins, giving the segments visible in figure 4.18 c). The number of projected points that falls inside each bin is recorded, and bin totals are normalised by the *median* bin value. In figure 4.18 c), the colourmap is centred on 1 such that segments in white are those in which the median number of points has been counted, red ones show ones where more than the median has been counted, and blue shows where fewer than the median have been counted. This figure reveals that there is a small amount preferential alignment of the orientations towards the cardinal directions (up, down, left, right, front, back), this is likely due to very small contact areas where the orientation can only be poorly defined.

A clear evolution in the distribution of contacts is visible: many contacts are gained in the x-direction and many are lost in the z-direction between configurations COEA01-GL-01 and COEA01-GL-17. Recalling that the sub-volume of grains on which this is calculated becomes part of the shear band that forms (whose normal is approximately $[-1, -1, 0]$), this seems to indicate that grain-to-grain contacts are lost in the direction of principal stress applied to the sample (along the axis of compression). Contacts are gained in the direction in which the shear band is advancing.

The work so far has been submitted to a conference on mathematical morphology – see Jaquet *et al.* (2013), and once developed further will be used to systematically analyse contact orientations in the images acquired in this work.




Figure 4.17: Equal area projection of contact orientations from 3D images of glass ballotini

Proposed solutions

While higher performance separation algorithms are tested, remedial measures have been explored starting from the contacts which are known to suffer from artefacts in their orientations. The approach has been to use the contact as a starting point, and extracting the grains in a way similar to Figure 4.7. The voxels representing the *surface* of the two grains near to the place at which they contact are incorporated into the measurement of the contact's orientation. This is done by intersecting a sphere (whose radius can be changed) centred on the centre of mass of the contact voxels and adding the voxels that represent the intersection of the sphere with the surface of the two grains in question to the calculation of the orientation of the contact. This addition of information helps disperse the contacts towards a more uniform distribution of contact orientations. The result for a sphere radius 6 pixels is shown in Figure 4.19, computed on the same contacts as those used for Figure 4.14 b), which are considerably improved – there are much fewer apparent artefacts in this distribution.

If further refinement is needed of this technique, the local perimeter of the grains could potentially each be fitted with a 3D ellipsoid, which may give a smooth normal for each grain, which could be averaged in order to get an overall contact normal.

4.2.6 Saving measurement information

From a technical point of view, these measurements of grain properties are saved in an expanded “Tabbed Separated Value” format in ASCII. This is a cleartext format (*i.e.*, which can be read by a text editor), in which information (such as label, 3D volume or position of centre of mass in the x -direction) pertaining to each grain is saved on a line, separating each field by a tab (ASCII




Figure 4.18: Equal area projection of contact orientations from 3D images of from a test on Caicos ooids (COEA01) at the beginning and at the end of a test




Figure 4.19: The orientations of contacts, calculated by taking into account the local perimeter of both grains

Character 09 or HT – Horizontal Tab).

The format is defined as expanded because each labelled object (in this case a grain) is

explicitly enumerated, and saved on a line whose number is equal to the label. If a single grain's information, whose label is 200 needs to be stored in a file, this would mean padding the file with 199 lines with zeros – except for the label column which will count up from 0 to 199. The 200th line would have the grain's information. It is clear that this expanded format is wasteful of memory, however compared to the memory requirements for the images on which these measurements the memory requirements are negligible.

This expanded format has been adopted for two reasons. The first is that this format allows the file to be read quickly into memory – where the same expanded structure is maintained. The fact that individual labelled objects (in this case grains) are each on a line whose number corresponds to the label of the object simplifies the reading into memory. A file describing 20 properties of 55000 grains occupies 12MB on disk and can be read in 1.5 seconds (several times faster than a spreadsheet program).

The second and principal advantage of this expanded data storage is the high speed operation that it allows. When dealing with grain-based processing, it is very useful to be able to refer to a grain directly (*i.e.*, by its line number). Suppose a grain label 200 is found in a labelled volume: a compact structure in memory, describing the grains requires searching through memory until grain 200 is found, which given the number of grains in the images becomes rapidly a bottleneck. The expanded structure in memory on the other hand, allows access to the grain's information on it's 200th row directly, this is very beneficial both in terms of speed, as well as simplicity of code.

This approach is maintained when dealing with contacts and other objects throughout the work.

4.3 Porosity (and void ratio) measurements

Porosity, or void ratio measurements are of fundamental importance in soil mechanics experimentation. They are used to characterise the state, in terms of density, of a granular assembly.

The relative density of a granular assembly has a profound effect on its mechanical response. Because of this, void ratio has often been used as a parameter in the modelling of the behaviour of soils: *e.g.*, in relatively advanced models like CamClay. These measurements represent, albeit in different ways, the solid-to-pore volume ratio:

$$e \text{ (Void Ratio)} = \frac{V_{Voids}}{V_{Solids}}$$

$$n \text{ (Porosity)} = \frac{V_{Voids}}{V_{Total}}$$

$$e = \frac{n}{1 - n} \quad \text{and} \quad n = \frac{e}{1 + e}$$

It is clear that these kind of measurements are of interest for this work.

In conventional soil mechanics testing, void ratio or porosity is a macroscopic measurement made on the entirety of a specimen. This is a useful measurement if the specimen is homogeneous, at least in its initial state. Conventionally, the volume of solids is considered not to change much, and the variation of the volume of the voids is calculated from the amount of water entering or leaving the specimen (in drained tests). As seen in Chapter 3, specimens are prepared dry, so the continuous measurement of the volume of voids is not possible in this manner. 3D images coming from x-ray tomography are however available at several stages during the test.

The pioneering work of Desrues *et al.* (1996) used x-ray tomography to study several traditional size triaxial specimens, *post-mortem* in a medical x-ray scanner. The voxels used in this study measured $700 \times 700 \times 4000 \mu\text{m}$ (the long dimension is aligned in the direction of the axis of the specimen). The Hostun RF grains studied (which have a reported mean grain size of $320 \mu\text{m}$) would not be individually visible in these x-ray scans. However the voxel size used is

small enough to contain only a few grains, a measurement which can be considered local, for the traditional size specimens approximately 100 mm in diameter (and 100 or 200 mm in height). With calibration of the measurements of the x-ray scanner, the reconstructed values of x-ray attenuation of each voxel can be related (knowing the density of the grains being scanned) to a void ratio or porosity. The ground breaking results of these scans show patterns of localised void ratio changes spread in complex geometrical patterns throughout the specimen.

This work, 16 years later, used a considerably more advanced (and suitable) x-ray scanner, allowing, as Figure 4.1 amply shows, the resolution of individual grains. It is possible to obtain a binarised image (with a good degree of confidence) of the solid and void phases of the specimen, in each state (see Section 4.1.1). This means that it is possible to calculate a macroscopic porosity or void ratio from each binarised image, inside the specimen’s domain in the image, the number of void and solid pixels can simple be counted, allowing porosity or void ratio to be calculated.

Local measurements of porosity, on the other hand, are not so straightforward. Since individual grains are present in the images, it is not useful to use each voxel as a measurement of void ratio or porosity (since each voxel will have “full” or “empty” value each time). Some coarse-graining or zooming out approach is required.

This section explores two potential methods for making local measurements of porosity. The first (Section 4.3.1) is a “continuum” style regular grid approach, where porosity is calculated on a small cubic subvolume at regular intervals. The second (Section 4.3.2) is a grain-based approach, which uses the granular skeleton as a reference system, and tries to associate a porosity to each grain, by tessellating the specimen space with tetrahedra joining grain centres.

4.3.1 “Continuum” porosity measurement

It is of obvious interest to be able to generate regular 3D porosity fields in order to allow a comparison with those reported in Desrues *et al.* (1996) (as well as many other since). In order to do this, the solid and void phases in the binarised 3D images coming from x-ray tomography need to be integrated over some subvolumes. The subvolumes need to be sufficiently small in order to make a usefully “local” measurement of porosity (*i.e.*, a subvolume occupying most of the specimen is not describing some local measurement of porosity). The subvolumes also need to be sufficiently large in order not to be *too* sensitive to the underlying material. This may seem counter-intuitive, but considering the extreme case of a subvolume of only one voxel, this would record porosities of 100% in the void phase and 0% in the solid phase – and nothing else: the result would simply be the binarised image that was used as input.

A subvolume size must be chosen so that when it is used to measure a granular material considered as having a homogeneous density, it provides a relatively stable measurement. When this can be said to be the case, then the subvolume is a Representative Elementary Volume (REV). Unfortunately there is an underlying problem that what is typically considered as a granular assembly of homogeneous density is not really very homogeneous at the grain scale.

Furthermore, it was considered by the author that a single REV may not necessarily be possible to define for an inhomogeneous granular material, such as one that has undergone shear localisation.

In order to study these two questions (appropriate size of REV and the possibility of different REV sizes in different parts of a specimen) some code was specifically written in Python.

The code, used for example, in Tagliaferri *et al.* (2011), uses on pre-binarised or trinarised 3D images (it has also been developed to work with images on unsaturated media – where not only solid fractions are important but the degree of saturation of the voids is also of great interest – see Kaddhour *et al.* (2013). In this study only binarised images are used). In these 3D image a set of nodes used for the analysis is defined (by default in an automatic and regularly-spaced fashion). Nodes are typically defined with respect to the image’s coordinate system, and so the nodes are an Eulerian reference frame relative to the moving material. For each node, the program grows a cube always approximately centred on the node. The cube grows step by step by shifting the

corners progressively away from the node (checking that the corners do not move outside the 3D image). As the cube grows, the porosity is calculated by summing the number of void pixels inside the cube and dividing this by the known cube volume. This generates (for each node) an array of cube dimensions and their corresponding porosity. These readings necessarily start from a volume of 1 voxel, which has either 0% porosity (solid voxel) or 100% porosity (void voxel). As the subvolume increases in volume, the calculated value of porosity changes. For nodes where the increasing subvolume does not go out of the domain of the specimen, the porosity reading is expected to tend towards the macroscopic porosity, for homogeneous specimens. For nodes where the subvolume eventually goes out of the domain of the specimen and into the void on the outside of the specimen, the value of porosity has a toward higher porosity (if the subvolume encounters the loading platens which have been binarised as solid, there will be a corresponding bias toward higher porosity). This is illustrated (for carefully chosen nodes) in Figure 4.20.




Figure 4.20: Evolution of the measurement of porosity as the cubic subvolume of on which it is calculated increases from one voxel, to a cube measuring $100 \times 100 \times 100$ voxels. This is shown, for specimen HNEA01, in state 01 for ten different nodes placed on the specimen’s axis and with 200 px vertical spacing. As the subvolume of calculation increases, the measured porosity can be seen to be tending towards a value representing the specimen’s overall porosity

In order to make measurements of porosity at a given node, a size of REV needs to be selected.

This can be done once for all nodes, by selecting an acceptable level of convergence using a graph like Figure 4.20. It is clear that a value of $10 \times 10 \times 10$ pixels is too small, and that $75 \times 75 \times 75$ pixels would be sufficient. However, the choice of a value between these two extremes is not easy. It is probably a relatively good idea to “play it dangerously” and prefer smaller values, since measuring with smaller subvolumes will make the measurement more sensitive to small features (at the expense of noise). In the example above, a fixed value of 28 has been chosen for the porosity calculations. The approach of tending towards smaller subvolumes also has the added bonus of rendering the calculation of porosity lighter from a computational point of view (many fewer data involved).

The program developed to calculate porosity also allows the size of REV to be automatically selected for each node. The program allows the input of a stability criterion based on the last x measurements of porosity as the cubic subvolume grows. This means that each node will be grown from a given subvolume size, and if (as an example of a stability criterion) the last three

readings are within 3% porosity of each other, then the the node is considered stable and the 3rd last reading of porosity is used. This allows measurements to be made on specimens where there are potentially different values of optimal REV size.

Unless explicitly stated, the first option is generally used in this work (to allow easier comparison with other work), but the availability of the second option means that the selection of the fixed REV size is rendered much easier with the easy and rapid generation of graphs such as 4.20.

The program has been written in such a way to allow easy modification of the shape of the REV (incurring a small performance penalty), an immediate shape that would be interesting to try would be a sphere. The effect of the shape of the REV however has not been considered in this work, and a cubic REV (except where trimmed when leaving the domain of the *image*) is used. Furthermore, the program is highly optimised; it takes advantage of the vectorisation of some parts of this problem, and therefore avoiding time-consuming loops (see 5.3.6 for an example of a vectorised problem), and can also run in parallel on multiple processors. For example, porosity can be evaluated on 1444000 nodes with an REV size of 28^3 on a $950 \times 950 \times 1600$ image in only 9 minutes and 9 seconds on a machine equipped with an Intel Xeon E5620 processor.

Results from this program are shown in Chapter 6. Results (*i.e.*, the values of porosity associated with the nodes spread through the 3D image) are presented as 3D images themselves, where each voxel represents a node. This is illustrated, in 2D in Figure 4.21

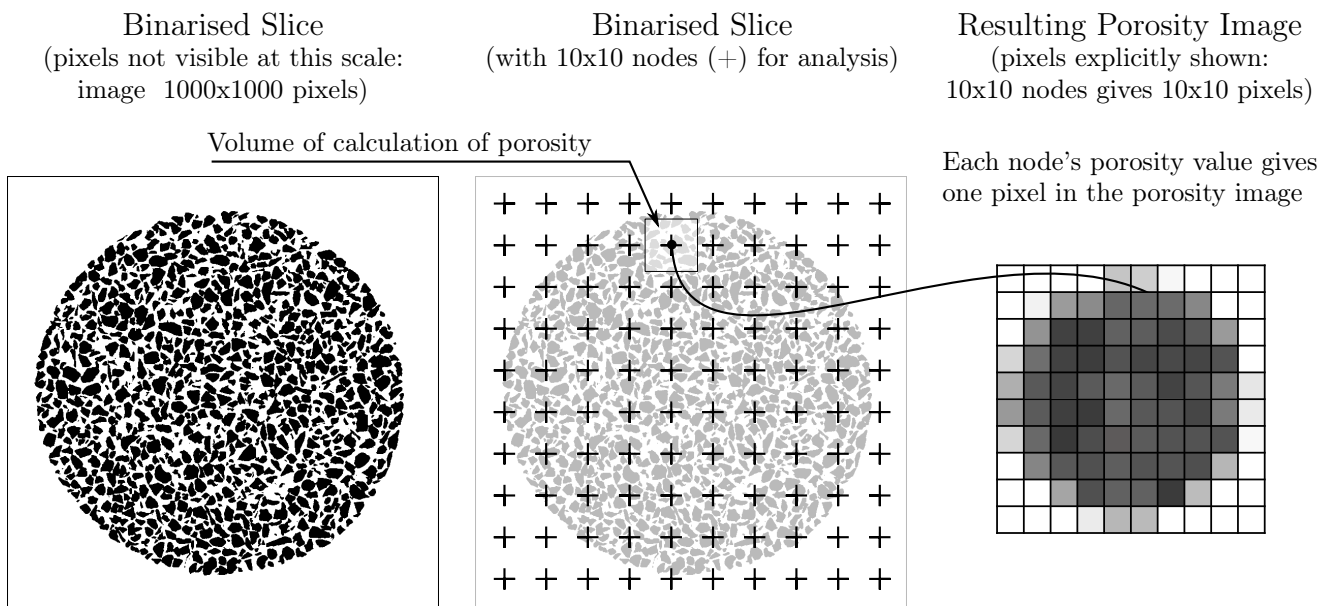


Figure 4.21: Diagram showing (in 2D) how each node, at which porosity is measured, corresponds to a voxel in a porosity image

4.3.2 Grain-based porosity measurements

Another promising method of measuring local variations of porosity, is to avoid the regular grid provided by the images, and instead to focus on the grains in the image. The cloud of points corresponding to the centres of mass of all the particles in a given configuration can be triangulated, so that the entire domain of the specimen is filled with many tetrahedra, each connecting four grain centres. The kind of triangulation used is the Delaunay triangulation (Delaunay, 1934), as computed by the Quickhull algorithm (Barber *et al.*, 1996) implemented in Qhull (with a SCIPY interface library called: “scipy.spatial.Delaunay”). For each grain centre,

the volume of all the tetrahedra having that grain centre as a vertex were thought of as an initial indication of the volume surrounding the particle. More correctly, Voronoi diagram (which is the dual graph of the Delaunay triangulation) should be used for this calculation, but in principle, both of these measurement should be sensitive to the changes in porosity at the grain scale.

The calculation of the volume around a grain was implemented, with each grain being assigned the volume of the tetrahedra having a vertex as that grain's centre and subtracting the grain's volume; results of this measurement were not able to reproduce continuum measurements of porosity, which are made with a high degree of confidence. The reason behind this failure is likely to be the nature of the triangulation used: in this case the triangulation algorithm simply connects nearest points in such a way that the entire domain is spanned. This means that some centres may have few tetrahedra connected whereas other may have more, simply based on some small changes in geometry.

A more objective measurement would either be to fix the number of nodes connected at each grain centre, or perhaps to connect grains based on their contact network in the style of Frenkel *et al.* (2008). These other types of triangulation are being investigated.

4.4 Measurement of specimen volume

A macroscopic measurement of the volume of the *specimen* can be useful for characterising the initial, macroscopic state of a specimen (such as through an initial porosity or void ratio – see Section 4.3). The ability to follow the evolution of these quantities during a test is important for the mechanical interpretation of each test.

The volume of the specimen is a subtly different thing to measure than the volume of the solid phase of the specimen (as done in 4.1.1). The volume of the specimen desired is the *macroscopic* volume, *i.e.*, the volume of matter (pore and solid) contained inside the specimen membrane and between the porous stones. Some simple tests show that it is not always possible to reliably identify the membrane in the 3D images coming from tomography.

This implementation therefore starts from the binarised image of solid and pore phases, with the objective of filling the pores inside the specimen.

A preliminary step is to remove the porous stones, which are often in the 3D image of the specimen, and which are of a density similar to the grains. This is typically done on the greyscale image of the specimen (where the distinction between the first and last grains of the specimen and the porous stones), which is then binarised. This step of porous stone removal from the images is relatively subjective, and can cause some standard errors in the measurement procedure.

A simple “fill-holes” algorithm is not suitable, since the grains that touch the membrane do not form a continuous boundary. The solid phase is therefore 3D-dilated 8 times (a number estimated to fill almost all of the pores in the specimens tested). A 2D fill-holes routine is then run slice-by-slice on the dilated binary volume. After this step, the solid phase is eroded 8 times, in order to obtain the original grain outlines, on the edge of the volume. The procedure is illustrated in 2D in Figure 4.22.

The result is a binary 3D image where voxels no longer represent solid and pore phases, but rather denote voxels inside the specimen and voxels outside the specimen. With the specimen volume 3D image thus created, the volume of the specimen can easily be computed by counting the specimen voxels.

The measurement of specimen volume, in order to characterise the isotropic compression behaviour of the main three granular materials studied in this work, is the object of Kate Wicher's internship at Laboratoire 3SR supervised by the author and Cino Viggiani – see: Wicher (2012). This work used the technique described above in order to measure changes in specimen volume – much smaller than the ones experienced by dense granular materials subjected to shear.

The characterisation of the propagation of error in this measurement technique has not been embarked upon, since the nonlinear process of filling holes is expected to be very challenging to

characterise satisfactorily.

4.5 Conclusions

This chapter has shown how the 3D greyscale imaged coming from tomographic scans on the dry granular materials tested can be analysed in order to yield separated and individually numbered grains, on which a variety of different quantities can be measured. The key measurements have been characterised to some extent. Furthermore, some tools for the specimen-level analysis of these images (in order to measure local porosity, or specimen volume) have also been developed and presented.

These data are crucial inputs to the following chapter 5, which will use these measurements – made on several images of different states – in order to measure grain or material kinematics, as well as the evolution of a variety of different measures described herein.

Furthermore, some of the measurements described are directly interesting for the static analysis of certain states, these are found in chapter 6, along with the findings made possible with the tools described in chapter 5.

Finally a table of the typical time required for each step is given in table 4.1.

	Calculation Time (minutes)
Calculating greyscale threshold	10 m
Binarising a 3D volume	3 m
Calculation of a “continuum” porosity	10 m
Separation with Visilog [®] 7.0	120 m
Calculation of grain properties with Visilog [®] 7.0	20 m
Labelling contacts	200 m

Table 4.1: Table showing approximate processing times for a single greyscale image coming from tomography

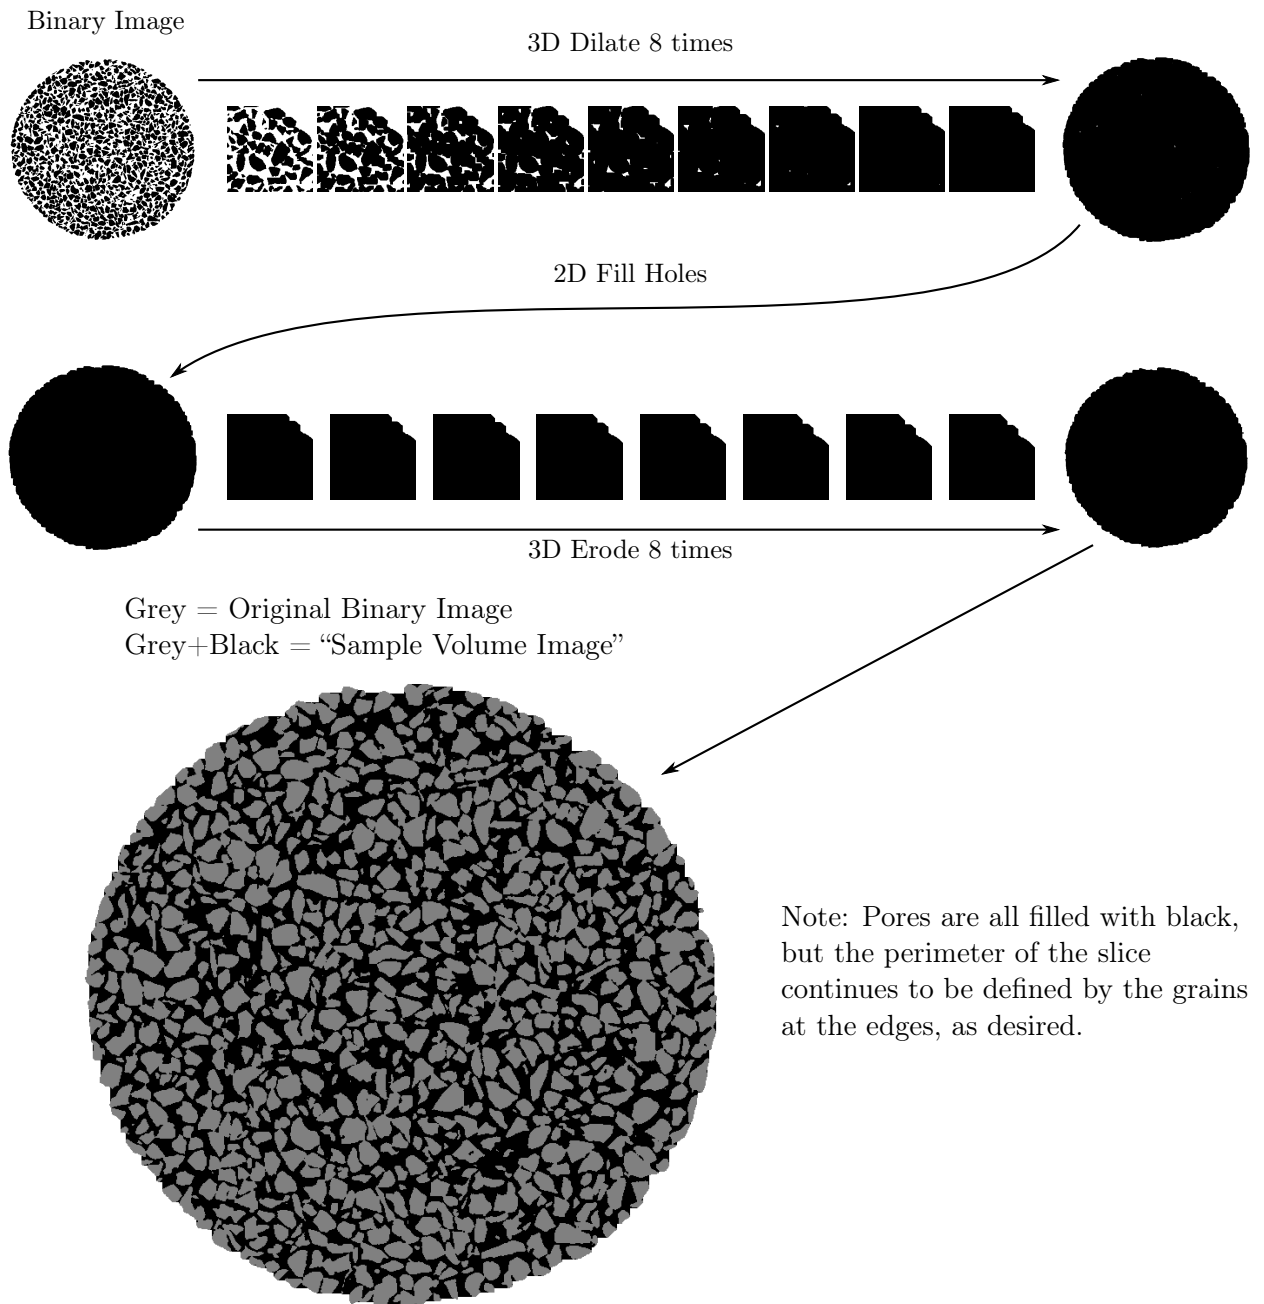


Figure 4.22: Illustration of the procedure used for calculating the volume of a specimen. The last image shows that the procedure, albeit simple works as desired – the perimeter of the shape is respected, and the inside of the shape is filled in

Chapter 5

Following processes in three dimensions

Many aspects of the response of granular media to load are well described by continuum models at a global scale. However, in reality (and in particular in the specimens tested herein), it is impossible to achieve a uniform distribution of state variables (such as the porosity), which means that, even in an initial state, it would be difficult to describe specimens as a homogeneous continuum. Furthermore, at the grain scale, sand grains are a very different concept from material points in a continuum. When a specimen undergoes strain, there will be a certain amount of tension between how individual sand grains actually respond, and what material points in the same positions would do in a continuum, due to the fact that grains will not be geometrically free to move in the way dictated by continuum mechanics.

Depending on one's philosophical position on continua, one could venture so far as to say that there is a de-coupling between what the sand *wants* to do (as a continuum) and what it *can* do (as a granular assembly). This de-coupling causes grain kinematics to become important, since they are non-affine. Macro-scale phenomena such as shear banding in sands may well find their origins in grain-scale mechanisms.

For the study of micro-mechanisms in a deforming granular medium, the ability to follow the material from one configuration to another is essential. Following the material's deformation between imaged configurations gives full-field measurements of the kinematics of all the material that composes the specimen being tested. The ability to compare these internal, 3D kinematics to the applied boundary conditions is the first step in the characterisation of the micro-mechanical response to globally imposed deformation.

The preceding chapter details the procedure for *imaging a material*. For example, a single 3D image coming from tomography can be processed in order to define the extent of the solid and void phases in the domain imaged by the x-rays. This in turn gives access to macroscopic measurements (such as specimen volume) and mesoscopic measurements (averaged over a Representative Elementary Volume), such as porosity. Further processing allows the definition of grains and contacts in the image which can be enumerated and measured individually, which gives access to many grain-based measurements such as coordination number, just to cite one. Such image processing can be performed on images of the same specimen in two different configurations, giving measurements in both configurations. It is of obvious interest to be able to follow the evolution of these measured properties, in order to see the interplay between micro- and macroscopic specimen behaviour. However, if the displacement of the material is not followed, then it is not possible to compare the evolution of measured quantities locally: these measurements can only be averaged over the specimen, and studied at that level (*i.e.*, by comparing the histogram of the coordination number of each grain in two different configurations).

If the material is followed through the different configurations in which it has been imaged, local analysis is possible. This can be thought of as *following a process*. Practically speaking, following the material's deformation requires the identification of invariant features in one (typically initial) configuration, and recognising these features in subsequent configurations. In this

kind of study, the most obvious invariant of features to study are grains (since grain breakage is not expected at the relatively low mean stresses experienced by the specimens). As Figure 4.1 has shown, grains are easily identifiable by eye in the 3D images coming from tomography, meaning that it should be possible to follow them between configurations. Different techniques for following the evolution of the material using images of grains have been developed and used, and are detailed in this chapter.

Once the material can be followed, and its kinematics measured, a truly local analysis of grain-scale mechanisms of interest is possible. Micro-mechanisms can be characterised either at the grain scale by looking at the properties of the grains involved (contact number, position, shape, size), or at a mesoscopic scale by relating the position of the grains involved to the porosity of the specimen in that position, for example.

Layout of chapter

This chapter shows how kinematic analysis can be carried out between different configurations (*i.e.*, between different images): Section 5.1 introduces Digital Image Correlation (DIC) as a tool for comparing and analysing images of deforming materials. This is followed, in Section 5.2, by the description of an implementation of perhaps the most natural DIC approach: Continuum DIC. TOMOWARP (Created by Stephen Hall and discussed in Hall *et al.* (2010), among others) is the Continuum DIC implementation used in this work. Continuum DIC follows (by 3D pattern matching) small (grain-sized, in this case) subvolumes of the images between configurations to give, for each subvolume, a 3D displacement vector mapping its displacement between the two configurations. Subvolumes are typically centred on a regular grid independent of the positions of the grains (which is why this technique is called “continuum”); the underlying discrete granular structure present in these images is ignored. However the regularly-spaced grid is convenient, a continuum mechanics measurement of strain comes naturally.

Section 5.3 then details grain-based kinematic measurements using ID-Track, developed in this doctoral work and published in Andò *et al.* (2012a), and developed further, with the addition of Discrete DIC (*i.e.*, grain-based DIC) in Andò *et al.* (2012b). This technique uses the grain-based measurements detailed in Section 4.2 and attempts to recognise each grain from a reference configuration by looking in the spatial neighbourhood of the deformed configuration for grains that match an invariant property (3D volume has mostly been used in this work). This technique simply aims to generate a mapping of grain labels between configurations. Once this is done, displacements of grains are easily computed by comparing the position of the centre of mass of the grain in each configuration. The measurement of rotations is more challenging, however, and two different techniques are detailed. The first technique is based on the comparison of a particle’s orientation in two different configurations (computed from the particle’s moment of inertia tensor). Since this is vulnerable to rather large errors, a second DIC-based technique has been developed, which is detailed in Section 5.3.6. Although grains are the natural reference frame on which to make kinematic measurements, the objectives of this chapter are to go further. This means using this granular reference frame to look at contacts (Section 5.4) or to look at neighbouring grains in order to obtain grain-scale continuum mechanics type measurements, such as strain (covered briefly in 5.5).

5.1 Principles of Digital Image Correlation (DIC)

Digital Image Correlation (DIC) is becoming an increasingly standard tool for the comparison of digital images coming from mechanical tests on materials. This kind of tool follows the displacement of “material points” from one image to the next. DIC comes from a heritage of non-digital image comparison known as “stereophotogrammetry”, which allows the comparison of two images with optical methods (see Butterfield *et al.*, 1970; Desrues and Viggiani, 2004, for

applications to geomechanics). With this kind of tool, the difference between the two 2D images is presented as 3D relief to the observer (to whom one image is shown for each eye) – this is the reason for which this technique is sometimes known as “False relief stereophotogrammetry”.

This section details the mathematical tool used to calculate the similarity between two images in this work – the Normalised Cross Correlation (Section 5.1.1). This is used to make rough measurements of the displacement between two images to the nearest pixel. Refinements – discussed in Section 5.1.2 – allow subpixel accuracy to be obtained in displacements, as well as the potential measurement of rotations (and any other affine transformation).

5.1.1 Cross correlation

With the advent of digital photography, the task of image comparison is shifted from a human operator’s eyes and into a computer algorithm. In order for this to work, an algorithm must have some tool that allows the similarity between two images, or parts of images, to be defined. A standard tool coming from the world of signals processing is the *cross-correlation* between two signals, which is a measure of the similarity between them. This is typically used for finding a small signal inside a larger one.

Normalised cross-correlation (NCC) was developed – specifically for the application to image comparison – in order to account for the fact that the same signal (*i.e.*, the image) could be stronger (*i.e.*, brighter) due to changes in illumination. Normalisation improves pattern-matching performance when lighting intensity is varying. Many other types of cross correlation have been developed (as well as other families of pattern matching such as the “Sum of Squared Differences”), but NCC is in common use in this domain and performs well. Pan *et al.* (2009) offer a comprehensive review of different algorithms used in 2D, which remains applicable in 3D.

The normalised cross-correlation between two 3D images (I_1 and I_2) is typically defined as follows:

$$NCC(u, v, w) = \frac{\sum_{x,y,z} I_1(x, y, z)I_2(x + u, y + v, z + w)}{\sqrt{\sum_{x,y,z} I_1(x, y, z)^2 \sum_{x,y,z} I_2(x + u, y + v, z + w)^2}} \quad (5.1)$$

Where:

- I_1 and I_2 are 3D images/matrices
- x and y and z are the coordinates of the images
- u and v and w are the integer displacement values to apply respectively to the x and y and z components of image I_2

One immediate remark about equation 5.1 is that u , v and w , the controlling variables of the *NCC* function, just shift I_2 around. Since I_2 is a matrix, and by definition discrete, in this version of the NCC calculation u , v and w must be integers, in order to make sense as indices to the matrix I_2 . A 3D image correlation is typically performed by scanning all the combinations of u , v and w in a given *search range*, and for each combination, calculating and recording the NCC. Once this is done (for a large search range this is a very repetitious and therefore slow process), the combination of u , v and w giving the highest NCC are selected as the best match between I_1 and I_2 . Since I_1 and I_2 are matrices representing pixelated images, the process described above is only able to measure displacements in the direction of u , v and w to the nearest pixel displacement, and is not able to measure rotations.

These two limitations (inability to measure rotations and displacements under the pixel) come from the fact that in the simplest sense, a signal is 1D, which means that rotational degrees of freedom, as there are in 2 or 3D do not exist, and so cannot explicitly be taken into account.

5.1.2 Refinements for subpixel and rotational degrees of freedom

As discussed above, the fact that the “signals” being treated are discrete 2 or 3D matrices means that equation 5.1 can only measure 3D displacements to the nearest pixel. In order to overcome this limitation, the discrete nature of the images needs to be transcended. There are two principal techniques to achieve this: through interpolation of the correlation coefficient or through interpolation of the images/signals.

Interpolation of the correlation coefficient

If only sub-pixel displacements are desired (*i.e.*, without taking rotations into account), the Normalised Correlation Coefficient can be interpolated. This is done by going back to the combination of u , v and w for which the highest NCC was obtained and fitting an $N+1$ Dimensional surface to the surrounding data in u, v, w, NCC space. The local NCC maximum of this smooth surface yields the coordinates of u , v and w with subpixel accuracy. The precision of this technique is a strong function of the quality of the fitting and the sharpness of the peak. This approach is used in TOMOWARP and has proved to be effective at yielding sub-pixel displacements. This approach, however, simply interpolates existing measures of similarity made during the scan of the displacement search range, and is therefore limited to simply giving higher precision in displacements.

Image interpolation

A more general, but more computationally expensive technique than the NCC interpolation approach is to interpolate the signals (I_1 or I_2). Interpolation of I_1 or I_2 allows the value of any position inside the image to be defined, based on surrounding pixels where image values are defined. This means that non-integer values of u , v and w could potentially be passed to the NCC function seen above. In this case I_2 would be calculated by interpolation. Suppose for example $u = 0.2$ and $v = 0$, I_2 in this case is constructed by adding 0.2 to the u value of each coordinate in I_2 , and performing an interpolation of I_2 in order to assign values to these coordinates. The interpolated I_2 is then used in the calculation of the Normalised Correlation Coefficient.

Furthermore, any affine transformation of the coordinates of I_2 is conceivable, since any coordinate can be interpolated. This means that in addition to translation, coordinates can also be sheared, rotated, scaled and/or squeezed. This obviously adds an enormous parameter space to explore, since a search range is required for each transform, and each comparison with I_1 requires the interpolation of all the coordinates of I_2 . This large parameter space is typically explored not in regular steps (since the number of interpolations and NCC calculations becomes rapidly enormous), but rather with an optimisation procedure. This is suggested, for example in Bay (2008), where a “Broyden–Fletcher–Goldfarb–Shanno” (BFGS) minimisation algorithm is used to find a minimum using the gradient of the NCC function with respect to all the variables used (for example $u, v, w, \theta, \phi, \psi$ for displacements and rotations.). The same author previously used a Levenberg–Marquardt algorithm for optimisation (Bay *et al.*, 1999).

If using a minimisation algorithm, the NCC, as shown in equation 5.1, needs to be rewritten so that the global minimum corresponds to the best match between I_1 and I_2 , this is typically done by optimising $(1 - NCC)$. As stated in Smith *et al.* (2002): “This “cross-correlation” function differs slightly from the traditional definition by a constant and a negation”.

5.1.3 Summary

The DIC tools described above allow mapping to be computed for a pair of images, I_1 and I_2 using the similarity criterion shown in equation 5.1. Typically the mapping of image I_2 onto I_1 is described by a single vector. In the simplest form described, this vector is composed of

integer values of displacement (in pixels). This can be improved by the implementation of the computationally efficient interpolation of the similarity criterion to yield sub-pixel measurements of displacements only.

A more general refinement that can be made is to interpolate I_2 , allowing the mapping from I_2 to I_1 to include any affine transformation (such as shearing, scaling or rotation of I_2) with relative ease. Each affine transform, such as a rotation around a given axis, can be explored and added to the vector describing the transformation from I_2 to I_1 .

Digital image correlation has been used both in a continuum sense in TOMOWARP (Section 5.2) as well as a grain-based implementation for the measurement of grain rotations as an addition to ID-Track, discussed in Section 5.3.

5.2 “Continuum” kinematic measurements

This section details the Continuum-DIC measurements on the greyscale images coming from x-ray scans. This technique is known as “continuum” simply because the kinematics of the material’s deformation are calculated at nodes (which can be thought of as material points) which are laid out on a regular grid. In the analysis of tomographic images of a clay rock for example, where (for the moment) the grain size is far smaller than the pixel size, this is the most straightforward option (see: Lenoir *et al.*, 2007, for an example of the application of Continuum-DIC to a deforming clay rock). The regular spacing of nodes is well-suited to tomographic images since the images provide a convenient grid (an other potential option is to base oneself on the boundaries of the specimen). In the current context, this approach is named a continuum one, because the discrete nature of the material being studied is purposefully ignored despite the fact that in the images acquired in the context of this work grains are visible by eye (Figure 4.1). The material is therefore being treated as a continuum. This in itself allows tools such as strain tensors from continuum mechanics to be easily calculated from the kinematics of the material point and its neighbours.

In this work the Continuum 3D Digital Image Correlation code TOMOWARP has been used in order to make continuum style measurements of the kinematics of the deforming specimen. The general principles of Continuum-DIC and TOMOWARP are outlined next.

5.2.1 Implementation

Generally speaking, the Digital Volume Correlation (Volume is used in this case to underline the 3D nature of the images), DVC, described by Bay *et al.* (1999) consists in the calculation of a mapping vector, for each node of the regularly spaced analysis grid. In this explanation, a 3D image of a reference configuration will be compared to a 3D image of a deformed configuration.

In TOMOWARP, the same nodes (*i.e.*, in the same places, with respect to the origin of the image), are defined in both the reference and deformed image. This can be thought of as an Eulerian specification of the position of the nodes, rather than a Lagrangian one, since they stay fixed with respect to the coordinate system of the image, rather than with respect to the material.

Keeping the same notation as equation 5.1 (where I_2 is the image being displaced by u , v and w), for each node a cubic subvolume called I_2 is extracted from the deformed image, and a cubic subvolume called I_1 of identical dimensions, and in the same position relative to the origin of the image, is extracted from the reference image. These subvolumes are centred on the node around which they are extracted. I_2 is generally known as the *correlation window*, and its size is typically an input to the correlation algorithm. The size of the correlation window is chosen in order to be big enough to capture the “invariant features” of an image, whilst being small enough not to capture too many invariant features, which may experience different kinematics between imaged configurations. In this case it is chosen to be roughly the size of one grain.

For each node, the search range is swept, displacing I_2 by integer values of u , v and w , and calculating, for each combination a Normalised Correlation Coefficient between I_2 ($x + u, y + v, z + w$) and the immobile reference image: I_1 (x, y, z). The result of this (time consuming) step is a set of u , v and w values which yield the highest NCC for each node. As discussed above (Section 5.1), two different strategies are available for the calculation of sub-pixel displacements. TOMOWARP implements the interpolation of the correlation coefficient, which is performed for each node after the pixel-resolution search. The mapping u, v, w from I_2 to I_1 , particularly where displacements are considered is typically considered as the displacement of a material point from one configuration to the next.

Very similarly to the illustration in Figure 4.21, the DIC procedure described performs measurements on regularly spaced nodes. As the figure shows, if nodes are regularly spaced, then a new 3D image where each pixel represents one node is a natural way of viewing this output. Since the output of this procedure, is for each node, a three-component mapping vector from one configuration to the next, three images are forthcoming: an image of all the u , v and w that gave the highest NCC at each node. These can be considered as 3D images, or fields of x , y and z displacement.

5.2.2 Strain calculation

The DIC technique described so far gives field measurements of material kinematics. These are able to reveal many interesting phenomena in these kind of images. However, for the study of deforming materials the ideal tool (for the analysis of this type of continuum measurement) is a field of the strain tensor, which can be derived from the kinematic fields.

Strain is a dimensionless quantity describing the relative movement of material points to which materials have a stress response – strain does not take into account rigid-body transformation that are inherent in the kinematics.

Thanks to the fact that, in this case, kinematics are calculated for nodes laid out on a regular grid, existing techniques for the calculation of strains can be used. In particular the techniques for the calculation of strain from 2D discrete fields of kinematics presented in Geers *et al.* (1996) is particularly applicable. TOMOWARP itself calculates strain for all the nodes of a specimen using a trilinear gradient function of the nodal displacements of cubic elements. This is output as individual 3D fields of the each component of the strain tensor, as well as the first and second invariants.

5.2.3 Output of data

Since TOMOWARP makes measurements based on a regularly spaced grid, scalar components coming from these measurements (such as the vertical displacement (Z) the XX component of the strain tensor, or even the maximum correlation coefficient achieved during the search process) can be saved as a 3D matrix – thanks to the regularity of the nodes – that describes the measured distribution of the scalar over the analysed domain. 3D matrices describing scalar quantities in space can conveniently be in turn considered as 3D images, and so all the image analysis tools that are used in this work, can also be used to visualise and measure the results of continuum image correlation.

5.2.4 Summary

Continuum-DIC is a very convenient tool for the measurement of displacements and therefore strains between imaged configurations of a geomaterial deforming. Thanks to the use of regularly spaced nodes where measurements are made, fields coming from this measurement technique are easily visualised and interpreted as images in themselves.

Since individual grains are visible in the images acquired in this work, a discrete, grain based tracking is preferred. However this Continuum-DIC is used in the analysis both to check

results, as well as to make intuitive measurements of strain in the specimen, since strain is easily calculated in this code from equally-spaced neighbourhood nodes.

5.3 Measurement of grain kinematics

At the relatively low mean stress at which these experiments have been performed, by far the most important mechanism of deformation is expected to be granular re-arrangement (grain breakage, for example, is not expected). The ability to follow all the grains that make up a specimen as the specimen deforms is of key interest for micro-mechanical analysis. This section describes how all the grains in a granular material can be followed between images by explicitly taking into account the discrete nature of the images acquired in this work. The technique developed in this doctoral work attempts to follow individual grains between configurations.

The preceding chapter (chapter 4) shows how individual grains can be defined and labelled in a given 3D image of a given granular configuration. The work reported in Hall *et al.* (2010) proposes a discrete grain-following approach that uses only a single labelled volume (segmenting very large volumes of data used to take a very long time). This technique, added to an existing program called CMV-3D (Bornert *et al.*, 2004) is briefly explained, but is not used in this work, since (thanks to a faster watershed algorithm) labelled volumes are available in all configurations.

To follow grains between *two* labelled images, each representing a different granular configuration, the label of each grain in a given (reference) configuration must be matched with the new label that that grain has in another (deformed) configuration. Once this apparently simple link is established between reference and deformed configurations, all the grain-based measurements detailed in Section 4.2 can be compared between configurations, meaning that each grain's changes can be followed from one configuration to the next. However, since the particles resulting from the segmentation procedure may not have a 1 to 1 correspondence to the grains being imaged (see Figures 4.4 and 4.5 for an illustration of the types of situation which may cause a non-physical segmentation), care must be taken to avoid errors in the establishment of the link between particle labels in reference and deformed configurations. If this can be done successfully, it lays the foundations for fully local grain-based analysis of the micro-mechanics of deforming sand.

Furthermore, once this link is established, kinematics of grains are a natural result. Section 5.3.3 shows how grain displacements are easily obtained from the difference in the position of the centre of mass of a given grain between configurations. The measurement of rotations is less straightforward and Section 5.3.4 details two techniques for the measurement of grain rotations. The first approach is based on orienting grains using their moment of inertia. This technique was deemed not to give results of sufficient quality and so was replaced by a grain-based DIC, which gives much more satisfactory results.

5.3.1 CMV-3D

The work that inspired this doctoral work, published in Hall *et al.* (2010), describes the analysis of a triaxial test on sand, both with the continuum technique TOMOWARP (described above), and with CMV-3D, developed by Michel Bornert and Yannick Pannier. CMV-3D is a 3D digital image correlation code that integrates processing libraries from ITK (see Yoo *et al.* (2002)), and is detailed in Bornert *et al.* (2004).

The discrete version of CMV-3D is a grain tracking code, that follows individual grains by a grain-based DIC approach. When this discrete technique was developed, the watershed algorithm being used was very slow (many days to treat a single volume). In part because of this, only a single labelled volume, describing a reference configuration is used in this implementation. Since only the greyscale image coming from tomography is available in the deformed configuration, only image correlation can be used to follow each grain.

Very briefly, for every particle in the reference configuration in the discrete-DIC approach in CMV-3D the voxels corresponding to the particle in the labelled volume are extracted, dilated and then used to mask the greyscale image of the reference configuration – as described in Figure 4.7. Digital image correlation, as described above in Section 5.1, was performed for each grain using the masked grain image as the correlation window. This was done very successfully in Hall *et al.* (2010).

There are two issues associated with the described approach, coming from the fact that only a single separated volume is used. The first is that, despite the high success rate in following grains, grain kinematics are all that it is possible to extract with this method: since configurations subsequent to the initial one are not segmented, this means extracting micro-structural information – such as details about contacts – is difficult for subsequent steps.

The second issue is that when segmentation errors occur in the creation of the labelled volume, these are handled unpredictably. In the case of two (or more) grains coming out as a single particle in the labelled volume, it is unlikely, given some deformation, that these grains do not move relative to each other. This means that it is quite possible that this particle is either lost, or has very unpredictable kinematics. In the other case, where a single grain is split into two particles this technique also has rather unpredictable behaviour: since the technique has proved to have a very high performance in finding grains, it also has high performance in finding parts of grains. It has been observed therefore that broken parts of grains are followed all the way through a test.

It is not immediately obvious whether the following of parts of grains is beneficial or not. For the calculation of a local strain field this is probably not too problematic, however for a very local analysis of contacts this likely causes problems.

This technique has been applied with a high degree of success in Hall *et al.* (2010) and although inspiration has been taken liberally from it, ID-track, the technique developed in this work (detailed below, in Section 5.3.2) differs on some philosophical points.

5.3.2 ID-Track: introduction

In order to make the link between labelled particles in a reference and a deformed configuration, a technique called ID-Track (Andò *et al.* (2012a), Andò *et al.* (2012b)) was developed. This approach was first suggested in Hall *et al.* (2008), but was developed from scratch in this doctoral work. The philosophy adopted in the development of this technique is one of simplicity and expandability. The technique was developed taking into account two main facts: the first is that, thanks to the labelling procedure, each grain has a unique label in every configuration. As the code extract 4.1 shows, it is easy to extract each particle and make measurements on the voxels that make it up – this information is part of the grain’s “identity card”. The kind of information that can be calculated is shown for two grains in Figure 5.1.

ID-Track operates on a matrix of data for each configuration, inside which the measured properties of each grain are listed for each grain (this data structure is briefly described in Section 4.2.6). ID-Track requires, at the very least, the position (typically the centre of mass) of each labelled particle, and at least one relatively invariant measurement of a grain property, with which a grain can be recognised.

The second principal idea behind ID-Track is that in a dense granular assembly, grains are not likely to move very far between images, which means that a specimen-wide search is not necessary for each grain. This can be implemented with varying degrees of sophistication, and in the work so far, has been implemented in the simplest way possible: grains are searched for in the vicinity of their centre of mass in the reference configuration. The vicinity is defined by a search range in 3D space, in a similar way to the search range for Continuum Digital Image Correlation, without the requirement for integer values of displacement. Improvements to this hypothesis could be made by assigning a probability to different possible granular rearrangements based on boundary conditions, and geometrical layout of grains. This is studied further in Section 5.3.8.

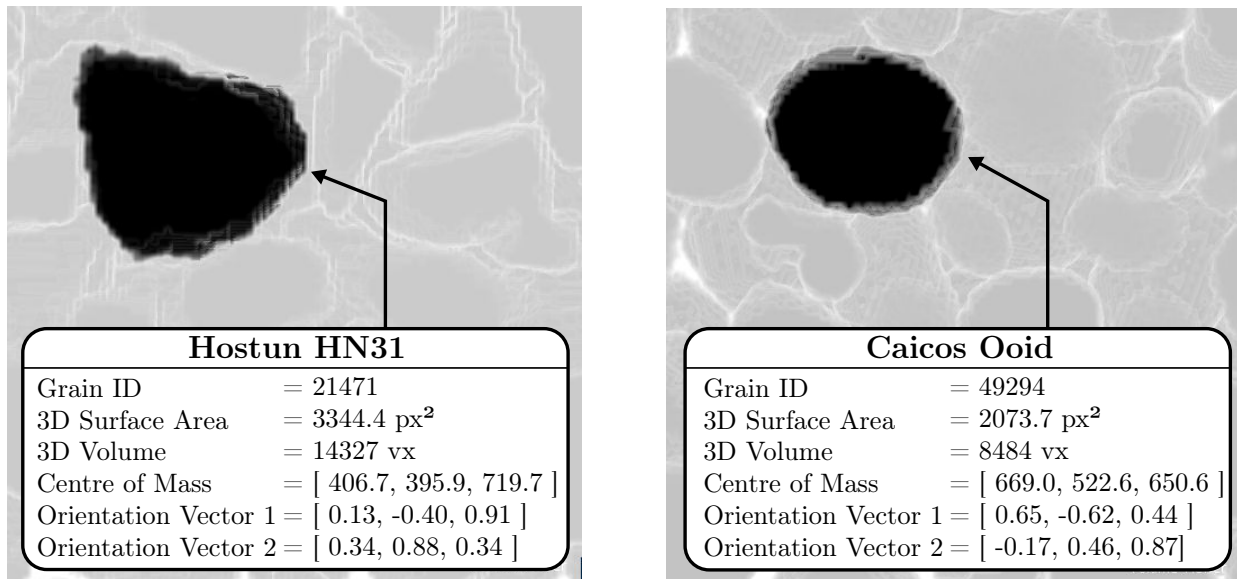


Figure 5.1: Figure showing two grains on which several different properties have been measured, in order to highlight the kind of “ID Card” which is available for each grain. Grains are visualised, along with their neighbours (faded out) in 3D. On the left a grain of Hostun sand is shown, on the right a grain of Caicos Ooid

Keeping in mind that each grain, in two different configurations has an ID, the flow chart of ID-Track is detailed in Figure 5.2. The flow chart shows how, for each labelled grain of interest in the reference configuration, all the grains in the deformed configuration whose centres of mass lie within the search window (which is relative to the centre of mass of the grain in the reference configuration) are selected (Figure 5.2 a).

Grain binning

The search for grains in the deformed configuration which are within a search range requires comparison of the centre of mass of each grain with the limits of search range, which can be a bottleneck when dealing with 50,000 or more grains. This problem is overcome by “binning” the data, *i.e.*, dividing up the domain of the 3D image into small, cubic subvolumes, which span the entire domain of the image, without overlap, and without leaving holes. Subvolumes of 50 pixels are used in this work, on average there are 7.6 grains per occupied subvolume. It is worth underlining that this procedure needs to be carried out only once, in the initialisation of the tracking – and not for each grain which is tracked. A list of the grains in each subvolume is maintained.

When looking for grains inside a search window, all the subvolumes which are intersected by the search window are determined (rapidly) from the limits of all of the subvolumes. All the grains present in these subvolumes (much fewer than in the whole volume) are then checked one by one against the limits of the search window. Thus many fewer grains are checked (those in the subvolumes in the search range *versus* all the grains in the volume), without ever incurring the penalty of a recursive search through all the grains in the entire volume.

Grain identification

Once a list of grains in the deformed configuration that lie inside the search range for a given grain in the reference configuration are obtained, some criterion is required to match the grain in the reference configuration to one of these grains. It no longer makes sense to use proximity in space, since this has already been used to defined which grains are potential matches. Furthermore, since




Figure 5.2: A flow chart describing the implementation of ID-Track. An X value of 10 is typically used in this work

the grains are expected to rearrange and move, the grain in the deformed configuration closest to the old position of the reference grain is not necessarily the same grain – especially since specimens undergo relatively large increments of macroscopic strain between imaged configurations.

To match the reference grain to a potential candidate in the deformed configuration, some invariant, measurable characteristics from the grain’s ID are needed. For example, the grain’s volume in voxels is a useful characteristic measurement, particularly since the tests being analysed are at a relatively low mean stress, with respect to the stiffness of individual grains: therefore only negligible amounts of volumetric strain are expected in the grains themselves. This means that throughout a test, a grain should have roughly the same volume.

A tolerance on this measurement is put in place to take into account the fact that it is possible for grains to change volume non-physically (as well as physically) as a result of the watershed process (a different amount of partial volume voxels may be deleted by a watershed separation line in each configuration, depending on the arrangement of contacts). For the measurement of volume a permissive tolerance of 10% is used. Figure 5.3 shows how the volumes of 9 grains evolve during a test. Despite some non-negligible changes of volume, it can be seen that these grains never have a change of more than 10% of their volume (the top and bottom limits of the graph).

As Figure 5.2 b) shows, all grains within a tolerance of a few percent are kept in consideration. This provides a list of grains that are inside the search range and inside a tolerance of the invariant feature (*e.g.*, volume) relative to the original grain in the reference configuration.

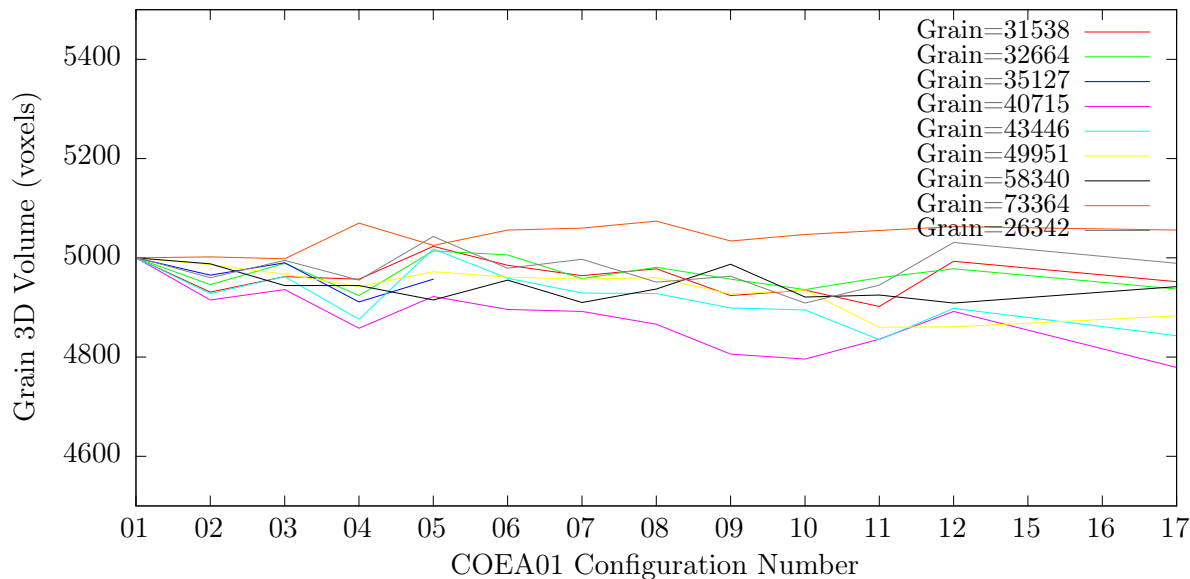


Figure 5.3: Evolution of the measured volume of 9 grains (all with the same initial volume) over the entirety of the test COEA01. Note how grain 35127 is not tracked all the way to the end of the test, and so only has information until COEA01-05.

Conflict resolution

Since there is often more than one grain that falls inside the search range and the tolerance, this must be handled by the algorithm. To resolve such “conflicts” the algorithm now uses two passes: on a first pass, the grain with the closest volume to the reference grain is temporarily selected as the best candidate. This first step is performed iteratively for every grain in the reference configuration, assigning each one a grain in the deformed configuration. If a grain in the deformed configuration, which has already been assigned to a grain in the reference configuration is found to be a better match for a subsequent reference grain, this new assignment is made, and the original reference grain is added back to the list of grains to track. This first pass, which simply matches grain volumes, is expected to work in most cases and give an approximately correct tracking for most grains.

The second pass involves using the median displacements resulting from the first tracking to calculate average displacements for each subvolume defined at the beginning of the procedure. The same procedure for finding grains in the deformed configuration for each reference grain is then followed. In the same manner as above, grains with an unvarying characteristic (such as volume) outside the tolerance of the reference grain’s characteristic, are discarded. The difference in the second pass is that the similarity criterion between grains includes not only the difference between grain volumes, but also the distance from the average displacement of the subvolume. These two criteria are weighted equally (although any weight is possible) and the grain in the deformed configuration that minimises the combination of both is selected. As before, this procedure is carried out for all the grains of interest in the reference configuration and any better matches are assigned, with the reference grain that is displaced being added back to the list of grains to track. The final output of this procedure is the mapping between grain labels from each reference grain to a grain in the deformed configuration, if a match has been made. If a reference grain has no matching grain in the deformed configuration, its corresponding deformed label is set to 0 (*i.e.*, the grain is simply not tracked in this given increment – it is *lost*).

Performance

ID-Track is able to correctly assign reference particle labels to deformed particle labels with a high degree of success, as long as a few conditions are met. The general condition for success is that false identifications are not made. The single unvarying characteristic used in this work – the grain’s volume in voxels – characterises each grain, but cannot be relied upon to be unique. As stated above this is for two principal reasons: first, the same grain may change volume (either due to a physical change of volume or due to the effect of separation), and second, in granular media (especially those with a tight grain size distribution) it is quite possible that at a given resolution two grains will have the same volume. A previous attempt to implement this technique on this kind of image was not successful because of this. With the large number of grains in the specimens studied, without implementing a local search window it is quite possible that between configurations there will be a different grain in the deformed configuration whose volume is closer to the correct equivalent grain from the reference configuration. By similar logic, if a search window is too large and includes too many grains with a volume too close to the reference grain being tracked, it is possible for incorrect assignment in the first pass of ID-Track. Enough of these errors could cause the median displacement of the subvolumes to be wrong, causing incorrect assignments in the second pass.

The problems (aggravated by having a tight grain size distribution) can be mitigated by reducing the amount of deformation between steps. This is often effective, since it means that a smaller search window can be used. In the images used these kind of tracking errors occur for much less than 1% of the grains being tracked, in other terms very rarely.

There is however, another source of error which can cause disturbances above the 1% of grains being tracked: incorrect separation of grains. More specifically, when there is an inconsistent separation of grains between configurations, this typically means that grains are not tracked. For example, if two grains are correctly separated into two particles in one configuration, and are stuck together in another configuration, the unvarying characteristic of volume would not allow either particle in the first configuration to be mapped to the stuck-together particle in the second configuration. This situation is detected (*i.e.*, no assignment would be made), but unfortunately leaves a hole in the data of the kinematics of the grains. Figure 5.4 shows vertical slices where grains that have not been tracked are highlighted. Table 5.1 shows the typical error rates in tracking due almost entirely to poor segmentation. In general about 4% of Hostun sand grains are not tracked, and about 1.4% of Caicos Ooid grains are not tracked. The number of untracked grains for Ottawa sand are relatively large, since the separation code has some trouble with the smaller grains in the specimen.

	% Grains not Tracked
Hostun	4.0%
Caicos Ooid	1.4%
Ottawa	6.2%

Table 5.1: Table showing % grains not tracked per increment for the different materials studied

Tracking an entire test

ID-Track has been used in this work to study successive increments: *i.e.*, tracking is performed between configurations N and N+1, as opposed to tracking each grain from the initial configuration. It was decided from the onset of this work not to re-relabel grain volumes (*i.e.*, once grain 1000 in configuration 01 is linked to grain 1033 in configuration 02 one could actually change 1033’s label to 1000 in configuration 02) once tracking is performed. The reason grains are not forcefully relabelled is that it seems like rather a final gesture, especially since there are still errors of a few percentage of grains being lost. Avoiding re-relabelling of grains also allows

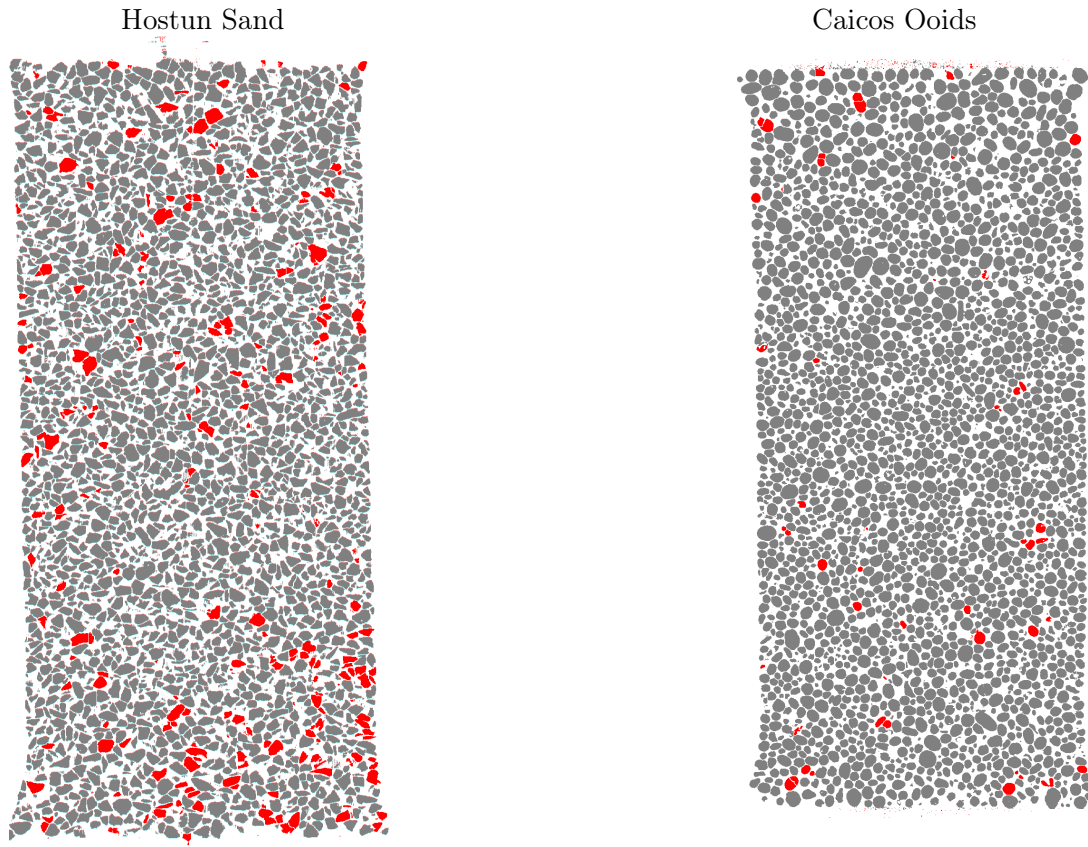


Figure 5.4: Slices through two specimens of Hostun Sand and Caicos Ooids, with grains that have been tracked shown in grey, and untracked grains highlighted in red

more flexibility in the re-processing of a given increment if particular errors are detected (such as a poor reconstruction for example). Once the different paths for the improvement of grain separation have been explored in order to give much-closer-to-perfect tracking, it is possible that a relabelling will be considered (which greatly facilitates the sharing of data, since grains have the same label in all configurations).

Regardless of the tracking strategy, if no re-relabelling is performed, each grain will have a series of different labels in all the configurations in which it is imaged. Therefore an algorithm has been developed to be able to follow individual grains through all said configurations. In its simplest form, this algorithm simply uses labels in the first configuration as the default labels for the grains, and follows each grain from the first configuration to the second directly by using the output of the ID-Track procedure. The tracking from the second configuration to the third is obviously done using the labels from the second configuration. The output of this tracking, is simply a list of the different labels (or identities) that each grain in the first configuration has in all subsequent configurations. A corresponding lookup table is also calculated and saved, allowing the lookup of any grain label in any configuration to the label of that grain in the first configuration.

If a grain is not tracked over any increment (due, for example, to an error in the grain separation process), it is only followed up to its the last configuration where it was tracked, thereafter the grain is lost. Over a single increment this is not very problematic, since the lost grains are only a few percent of the total number of grains. Over many increments, grains are increasingly lost, which when cumulated over a test can cause problems in the interpretation of the data. In order to resolve this, for each configuration after the first one, the algorithm checks for grains that start to be tracked from this slice, and relabels them as “foundling grains” and

assigns them a fictitious grain label (bigger than biggest grain label in the first configuration). No attempt has been made (for the moment) to relate these foundling grains to grains in the original configuration, although this is discussed as a potential improvement to the technique in Section 5.3.8.

5.3.3 ID-Track: measurement of displacements

Once a reliable link between labels in different configurations exists, it is relatively simple to calculate the displacement of each tracked particle over each increment. Each particle's label in each configuration is used to look up the particle's centre of mass – represented by a three-component vector – in each configuration. The difference between these vectors gives the displacement of the grain's centre of mass between configurations.

The calculation of the grain's centre of mass dictates the metrological aspects of this measurement – this has been discussed and characterised to some extent in Section 4.2.1. In the measurement of displacement, the error associated with the measurement of position is simply incurred twice. As stated, the maximum error, in each component of the measured displacement is expected to be less than (but in the order of) 0.05 pixels, which corresponds to 0.75 μm for each component, in the case of these scans. This means that the maximum expected error in the measurement of displacement is expected to be less than 1.5 μm per component.

5.3.4 ID-Track: measurement of rotations

Grain rotations have proved to be key measurement related to the formation and development of shear bands. In 2D, rotations of grains are relatively easy to picture, since a single rotation angle (a simple scalar) fully defines the rigid-body rotation of a grain. Conveniently, the sense of rotation (clockwise or anticlockwise) is a natural and intuitive quantity. In three dimensions rotations are conceptually more challenging, and have thus been approached from different directions, giving a rather diverse set of tools for thinking about and dealing with rotations in 3D. Rotations in 3D turn out to be very important in 3D graphical applications, such as animation and computer games, which means that tools and literature for dealing with rotations in 3D in software are readily available. This section begins by detailing some of the different systems for defining 3D rotations. The two different techniques for measuring 3D rotations of grains developed in this work are then detailed. The first technique, based on grain orientations coming from the moment of inertia tensor, is vulnerable to some rather large errors in the measurement of rotation, and so has been supplanted by an improved, grain-based DIC technique.

The two techniques are compared in the following section (5.3.7), on images where a known rotation is imposed on a collection of grains. Furthermore Section 5.3.10 contains some useful mathematical tools for the management of 3D rotations.

Systems for definition of 3D rotations

Rotations in mathematics are a subset of affine transformations. Rotations describe the motion of objects around a fixed point. This fixed point – the centre of rotation – is typically the origin of the space.

The most natural mathematical description of a rotation is an $N \times N$ matrix (where N is the number of dimensions of space being worked in). This matrix can be used to multiply vectors representing points in the space, in order to rotate them around the origin.

Three Euler angles are also sometimes used to describe 3D rotations. Euler angles rotate a reference frame, and so when three angles are used to describe a rotation, each angle rotates the entire reference frame. This means that the order in which the angles are applied is important. Euler angles are subject to three main problems: the first is that there are many different standards for the order in which the angles are applied – which is a source of confusion. The

second problem is a technical one: Euler angles are vulnerable to a phenomenon known as “Gimbal lock” – this is where a rotation can rotate the axes of rotations of the coordinate system so that two of the rotation axes are identical, which means that only two-dimensional rotations can be explored. Since the objective is to measure 3D rotations of grains, if the system describing the rotations is able to get locked in this way, it is possible that the desired rotation is not accessible. The last problem with using Euler angles is that they are not very easy to interpret independently – since three different rotations need to be applied in a specific order.

3D rotations can also be expressed as a rotation around a given axis. This is a convenient and intuitive framework to work in since there are many fewer scalars to keep in mind. In this work, this framework is used albeit with a popular modification, whereby the axis and angle are described together by a single *rotation vector*. Instead of giving a unit vector describing the axis of rotation, and a scalar describing the angle of rotation, these two are combined into a rotation vector whose direction is the axis of rotation, and whose norm is the angle of rotation. This compact notation is convenient not only for the interpretation of results, but also allows efficient storage on disk. The Rodrigues’ rotation formula (derived in Koks, 2006), and illustrated in equation 5.10 allows the calculation of a rotated vector from a rotation expressed in axis and angle form.

Quaternions are hypercomplex numbers, which can also be used to describe rotations. Quaternions are a powerful way of describing and manipulating rotations, which are very popular in many applications. Quaternions are made of four scalars, which are most easily defined starting from the axis and angle representation of a 3D rotation. Applying a 3D rotation described by a quaternion can be done in a very efficient manner.

It is also relatively easy to convert between different types of representation of 3D rotations: formulas exist to and from all types of rotation to the rotation matrix representation.

Summary

The axis and angle representation, using a rotation vector is used in this work. This will be represented by three scalar components (and not angles): θ, ϕ, ψ . This has been chosen principally for reasons of intelligibility of results, as well as efficient storage. The Rodrigues’ rotation formula is very convenient for rotating a single vector, but when many coordinates need to be rotated it requires too many operations on each coordinate. In these cases the 3D rotation is converted to a rotation matrix and this is applied by matrix multiplication to all the coordinates at the same time.

5.3.5 ID-Track: measurement of rotations from grain orientation

As discussed above, the measurement of the displacement of each grain is obtained by calculating the difference in the position of its centre of mass in 3D from one configuration to the other. A natural approach, and historically the first approach attempted in this work, is to perform the same kind of calculation for the measurement of rotations. This is to say that a characteristic 3D *orientation* is calculated for a grain in a each configuration and the comparison of this orientation between configurations gives a 3D rotation.

Different methods exist for measuring orientations of particles, the previous chapter details two potential options: the “feret” orientations (Section 4.2.4) of the width and length of the particle could be turned into an orthonormal base, otherwise the maximum and minimum eigenvectors of the moment of inertia tensor (Section 4.2.2) can also be used to orient the particle. The latter has been adopted for this technique. The maximum and minimum eigenvectors of the moment of inertia of a 3D particle are two characteristic vectors that point in the grain’s longest and shortest directions. These are calculated on the labelled images of the particle in two different configurations, and then an orthonormal base is constructed in each configuration. The orthonormal base in a given configuration is calculated by normalising the long axis (vector

“length”, l) and using it as a starting point. The short vector’s normal component is used as the second vector in the base (vector “depth”, d), and a third vector normal to both (vector “width”, w) is then calculated in order to complete the orthonormal base. The rotation of this base, between configurations (called reference and deformed configurations) is then given by constructing the rotation matrix for this transformation:

$$R = l_{def} \cdot (l_{ref})^T + w_{def} \cdot (w_{ref})^T + d_{def} \cdot (d_{ref})^T \quad (5.2)$$

The *Null Space* or *Kernel* of the rotation matrix is Calculated:

$$a = Null(R) \quad (5.3)$$

such that $R \cdot a = 0$ (a should be a 3-component vector representing the axis of rotation)

The angle of rotation is given by rotating a vector normal to the axis of rotation and measuring the rotation that it undergoes. The normal component of any vector is acceptable, a unit vector in the z direction z is used.

$$n = z - \frac{a \cdot z}{a \cdot a} \cdot a \quad (5.4)$$

a normalised n (\hat{n}) is then rotated to give the angle of rotation (α):

$$\hat{n}' = R \cdot \hat{n} \quad (5.5)$$

$$\alpha = \arccos(\hat{n}' \cdot \hat{n}) \quad (5.6)$$

This technique for the measurement of rotations, whose performance is evaluated on real data in Section 5.3.7, is vulnerable to the artefacts discussed in Section 4.2.2, which introduce non-negligible errors for the grains studied in this work. These errors can be drastic when there is some symmetry in the long or short directions of the grains.

One great advantage of the described technique is that it involves very few computations. The moment of inertia tensor can be obtained for a 3D object with relative ease, by simple summations as described, for example in Novotni and Klein (2001). Eigenvectors can be extracted with very optimised algorithms, and thereafter the calculation of the 3D rotation shown above is also not computationally intensive. In short, in cases where the orientations of grains is well characterised by the moment of inertia tensor, this is a good option for the measurement of particle rotations. Unfortunately in the case of the particles studied in this work the moment of inertia tensor is not sufficient to describe a uniquely orientation.

5.3.6 ID-Track: measurement of rotations from grain-based 3D DIC for rotations

When grains cannot be oriented reliably in 3D images by three orthonormal vectors, the measurement of rotations needs to be enriched by bringing more data into the calculation. As shown in Hall *et al.* (2010), grain-based DIC appears to work well on these kinds of images.

As discussed in Section 5.1, there is only one approach to image correlation that allows the measurement of rotations by allowing the rotation of one of the images being correlated: subpixel

refinement through image interpolation. Figure 5.5 gives an outline of how grain-based DIC has been implemented in this work for the measurement of the rigid-body rotations of grains.

Techniques that use image interpolation for image correlation are very computationally expensive. Therefore to speed up the entire process ID-Track is used to determine grain displacements and to match grains before performing the discrete DIC. This avoids the initial, slow process of finding the initial displacements by image correlation. For speed, grains are not searched for over a large displacement search window using image correlation (as was done in: Hall *et al.*, 2010). Instead in this implementation, the grain tracking provided by ID-Track (which maps grain labels from a reference configuration to a deformed one) is used to isolate the grains in question, so that rotations can be measured.

Extraction of grains

Labelled volumes exist for both configurations which means that the labelled voxels of a given grain can easily be extracted from both volumes – making sure to use the correct label (from the tracking process) in the deformed configuration.

A PYTHON code example showing how a labelled object can be extracted (Code extract 4.1) is available in the previous chapter. This shows that relatively optimised functions already exist for the identification of the extents – or “bounding box” – that a given labelled particle has in a labelled image. The location and extent of this bounding box can then be used to extract the grain’s labelled voxels in each configuration. All labelled voxels belonging to other grains are then deleted from each subvolume – this corresponds to step 1 in Figure 5.5.

At this stage of the process, a 3D subvolume containing only the labelled voxels of a given, tracked grain exists in each configuration. There is no guarantee that these two volumes are the same shape – the tightest bounding box may be different for a given grain in the two configurations under consideration, particularly if it has rotated. The next step is to copy each grain into a cubic subvolume of fixed size. There are several reasons why identically shaped cubic subvolumes are interesting for this process, but by far the largest is code simplicity – especially in the later stages of image correlation. In order to achieve this, the largest of all three dimensions of the bounding box in both configurations is defined as l . The dimensions selected for the cubic arrays for both the reference and deformed configurations, is such that a cubic grain being rotated around the centre of the array does not go outside the array. Furthermore, the cubic array is padded by a certain number of pixels in order to allow dilation of the labelled grain voxels (see Equation 5.7).

Length of one edge of the cubic array for each grain is given by:

$$NewDimension(\text{rounded up to nearest integer}) = \sqrt{3 \times l^2 + (2 \times padding)} \quad (5.7)$$

The dilation of the extracted grain images is required when the labelled grain voxels are used as a mask to fetch the greyscale information of the grain: dilation of the grain allows the greyscale of the boundary of the grain to be considered, as explained in Figure 4.7, in the previous chapter, as was done in Hall *et al.* (2010).

Two new arrays of this dimension are allocated, for the reference and deformed configurations. The centre of mass of the labelled grain voxels is then calculated and rounded to the nearest pixel. This is used to place the labelled voxels of the grain approximately (to the nearest pixel) in the centre of each cubic array. This operation is done to the nearest pixel, in order to avoid interpolation – both for speed as well as to avoid the degradation of data by multiple interpolations – this completes step 2 in Figure 5.5.

At this stage both grains in both matrices are dilated $padding$ times (step 3a in Figure 5.5). In this work, $padding$ is typically set to 1 – but this has been kept a variable, since a value of




Figure 5.5: Flow chart showing the implementation of the 3D grain-based image correlation for ID-Track, allowing the measurement of rigid-body rotations and translations of grains. N is equal to one in this work

3 is reported in Hall *et al.* (2010). A value of 1 allows the darker partial volume voxels to be captured, while minimising the possibility of capturing pixels that belong to other grains – as shown in Figure 4.7.

The greyscale information for each grain is collected from 16-bit greyscale versions of the reconstructed greyscale volumes. Since these are the same size as the labelled volume, the bounding boxes calculated in the first step are combined with *padding* and the displacement required to centre the grain, to fetch the correct greyscale voxels in each case. The dilated, labelled grain voxels (3b in Figure 5.5) are used as a mask for the greyscale data: the greyscale voxel values are only copied into the cubic arrays where there are dilated grain voxels, in other positions the values remaining zero. At this point, two centred, greyscale grains, are ready for comparison in two different, but identically sized cubic arrays.

Interpolation

Since any arbitrary grain rotation is possible, it is expected that many different 3D rotations will have to be tested before getting a match between reference and deformed configurations. This means the generation of many interpolated images before a maximum Normalised Correlation Coefficient is reached. This is certainly the most iterative part of this procedure for the measurement of 3D rotations, and is consequently expected to be the slowest. Some pre-processing is thus done in order to minimise the number of repeated calculations in the search for the Normalised Correlation Coefficient.

Trilinear interpolation has been used (linear interpolation in 3D) in this work and a technique has been developed (detailed below) for the pre-processing of images to accelerate interpolation (which is in principle applicable to more sophisticated interpolation schemes).

Very simply, interpolation of a 3D image consists in trying to reconstruct a greyscale value for a coordinate in the image domain that does not fall directly on a pixel where there is information. To give a 1D example, if a pixelated line is considered, with values defined at every pixel, if a value is needed at pixel position 7.2, this must be reconstructed, typically from values of the immediately neighbouring pixels (7 and 8). If in this case linear interpolation is used, if pixel 7 has a value of 0 whilst pixel 8 has a value of 100, a linear interpolation would give 20 for pixel position 7.2.




Figure 5.6: A cube, representing 8 voxels (each vertex of the cube is a voxel, and contains a discrete grey value). The definition of a value for point *A* which is not a vertex requires interpolation

In 3D this means that 8 neighbouring points must be taken into account. Figure 5.6 shows a unit cube, centred on the origin which contains a point *A* (coordinates *x*, *y* and *z* which are all between 0 and 1). Each vertex of the cube represents a voxel, whose value is known. The formula for trilinear interpolation (available for example in Hill, 1994) gives the following for point *A*:

$$\begin{aligned}
A = & V_{000} \times (1-x) && \times(1-y) && \times(1-z) && + \\
& V_{100} \times (x) && \times(1-y) && \times(1-z) && + \\
& V_{010} \times (1-x) && \times(y) && \times(1-z) && + \\
& V_{110} \times (x) && \times(y) && \times(1-z) && + \\
& V_{001} \times (1-x) && \times(1-y) && \times(z) && + \\
& V_{101} \times (x) && \times(1-y) && \times(z) && + \\
& V_{011} \times (1-x) && \times(y) && \times(z) && + \\
& V_{111} \times (x) && \times(y) && \times(z) && +
\end{aligned} \tag{5.8}$$

where V_{abc} is the Value of greyscale at point a, b, c .
and $0 \leq \{x, y, z\} \leq 1$ and $\{a, b, c\} = \{0, 1\}$

Consider now a 3D array of coordinates, containing elements such as A at which a given image should be evaluated. If the coordinates in this array do not fall exactly on integer coordinates where values in the image are defined, the values for the points must be interpolated, using for example trilinear interpolation as defined above in equation 5.8. It is possible to write a loop that interpolates each element (*i.e.*, each point) of the array sequentially, but this rapidly becomes a substantial bottleneck when used in this context, which requires several hundred different interpolations to be made *for each grain*. The interpolation of an array of coordinates can however be *vectorised* so that instead of writing a loop over each element of an array of coordinates, all the coordinates are calculated at the same time. For example, this can be performed taking advantage of PYTHON's (or any other language for scientific computing) optimised handling of matrix operations.

To take advantage of vectorisation, equation 5.8 – which is only valid for points inside a unit cube starting from the origin – must be applicable. Thankfully this is not hard to do; code extract 5.1 shows how this can be done in relatively few lines of PYTHON, largely taking advantage of matrix manipulation. The crucial point of this implementation shown is that an array of any dimension is handled without loops, and in a single run of the function.

This speed boost can be further augmented by taking into account the fact that in this implementation of interpolation for image correlation, the same image is going to be interpolated repeatedly. The second last line of the program takes by far the longest amount of time, since the different values of “image” have to be looked up $3 \times 8 \times$ Number of Coordinates times, which requires a lot of operations on memory. The multiplication of the number of memory lookups by 8 is because values of the image are fetched in the eight corners around any given point (see Figure 5.6). These repeated lookups can however be avoided at the cost of increased memory consumption.

The solution used to recode the number of operations in memory in this work aims to avoid the individual multiplications of the array of coefficients also inherent in the second last line of the code extract. In order to do this, the 3D image to be interpolated is expanded into a fourth dimension so that each point no longer contains a simple scalar but instead contains a $2 \times 2 \times 2$ cube of scalars (in the style of Figure 5.6, with the original scalar at 0, 0, 0 and the neighbouring points in the other corners). This neighbourhood of greyscale values is embedded inside each point. This obviously is very wasteful of memory (which is used very redundantly), however the pre-computation of this image means that local values around a given point can be accessed in a single place in memory.

The creation of an $8 \times N \times N \times N$ image takes a certain amount of time – and may not save time if only a single interpolation is required. However, since many interpolations are expected,

Code 5.1: Specimen Python code for a trilinear interpolation function showing preparation of an array of 3D coordinates for use of equation 5.8

```

# Definition of a trilinear interpolation function taking a 3D image as input ,
# and an array of coordinates at which the 3D image it should be interpolated.
# N.B. when coordinates are not integers the image must be interpolated.
# Inspired by http://stackoverflow.com/questions/6427276/3d-interpolation-of-numpy-arrays-without-sciPy
def trilinear_interpolation( image, coords ):
    # Generate an array of coordinates which are rounded down,
    # (casting to int in python truncates)
    coords_0 = coords.astype(numpy.integer)

    # Now add 1 to the rounded down array. This makes a array of coordinates which
    # is a rounded up version of the original coordinates array.
    coords_1 = coords_0 + 1

    # Also generate a array of coordinates in a local reference frame (of the voxel),
    # i.e. subtract the rounded down coordinate array from the original array.
    # The coordinates in this array are all between or equal to 0 and 1.
    coords_rel = coords - coords_0

    # A last array of "1 minus" the local array of coordinates is generated.
    # This too must have values between 0 and 1 inclusive.
    one_minus_coords = 1 - coords_rel

    # An array of coefficients, 8 in total for each coordinate is initialised.
    # coords.shape[1] is the number of coordinates.
    all_coeffs = numpy.zeros( (8, coords.shape[1]) )

    # There should be a step here to make sure that the coordinates
    # do not go outside the image, which would cause an error.

    # Now the array of values for each coordinate can be calculated in two steps.
    # The first is to calculate the 8 coefficients that multiply each component
    # such as V000.
    all_coeffs[0] = (one_minus_coords[0])*(one_minus_coords[1])*(one_minus_coords[2])
    all_coeffs[1] = (coords_rel[0]) *(one_minus_coords[1])*(one_minus_coords[2])
    all_coeffs[2] = (one_minus_coords[0])* (coords_rel[1]) *(one_minus_coords[2])
    all_coeffs[3] = (one_minus_coords[0])*(one_minus_coords[1])* (coords_rel[2])
    all_coeffs[4] = (coords_rel[0]) *(one_minus_coords[1])* (coords_rel[2])
    all_coeffs[5] = (one_minus_coords[0])* (coords_rel[1]) *(coords_rel[2])
    all_coeffs[6] = (coords_rel[0]) * (coords_rel[1]) *(one_minus_coords[2])
    all_coeffs[7] = (coords_rel[0]) * (coords_rel[1]) * (coords_rel[2])

    # In this second step those coefficients are applied to the greyscale values
    # in the image. Note that coords_0 and coords_1 here are arrays of coordinates,
    # therefore coords_0[0], for example, is a array of x-coordinates.
    # Python allows this array indexing.
    interp = ( image[ coords_0[0], coords_0[1], coords_0[2] ]*all_coeffs[0] +
               image[ coords_1[0], coords_0[1], coords_0[2] ]*all_coeffs[1] +
               image[ coords_0[0], coords_1[1], coords_0[2] ]*all_coeffs[2] +
               image[ coords_0[0], coords_0[1], coords_1[2] ]*all_coeffs[3] +
               image[ coords_1[0], coords_0[1], coords_1[2] ]*all_coeffs[4] +
               image[ coords_0[0], coords_1[1], coords_1[2] ]*all_coeffs[5] +
               image[ coords_1[0], coords_1[1], coords_0[2] ]*all_coeffs[6] +
               image[ coords_1[0], coords_1[1], coords_1[2] ]*all_coeffs[7] )

    # Return the interpolated values of the image at the coordinates which
    # were passed to the function as input.
    return interp

```

Code 5.2: Specimen Python code for an alternative version of the last line of code example 5.1 when an $8 \times N \times N \times N$ image is used as an input. Indentation changed in order to fit equation on one line.

```
interp = ( image[ coords_0[0], coords_0[1], coords_0[2] ].T * all_coeffs ).sum( axis=0 )
return interp
```

this pre-computed matrix can be saved and re-used for all interpolations of a given grain. Code extract 5.2 shows the new version of the last two lines of the trilinear interpolator shown in previous extract (5.1) when an $8 \times N \times N \times N$ image is used as an input. The “sum(axis=0)” makes a sum across all eight coefficients (as per equation 5.8) after they have been multiplied by the respective greyscale neighbourhoods in the image, coming from the fourth dimension of the image, which is now only addressed once.

This relatively optimised tool allows many 3D interpolations on an array of coordinates to be made relatively quickly (*i.e.*, faster than the equivalent library in SCIPY for the data involved with this work). A simple evaluation of the speed of the different types of trilinear interpolations discussed is shown in table 5.2, where different numbers of iterations are performed on a $50 \times 50 \times 50$ cubic image, on an Intel i7-640LM processor. Row (a) in Table 5.2 shows the time required to interpolate the $50 \times 50 \times 50 = 125000$ points in the image in a loop that interpolates each point individually. In both relative and absolute terms this is very slow. The speed of interpolation is substantially improved by vectorising the problem, so that all the array of coordinates is interpolated at the same time. Vectorisation gives a great improvement in speed (Table 5.2 row b).

Table 5.2 row d) shows the further speed boost obtained when embedding neighbourhood greyvalues inside each point. Doing this avoids the memory operations involved in collecting 8 neighbouring values for each point. For a single interpolation however, the setup time takes approximately as much time as is saved by avoiding the extraneous operations in memory. Consequently, the two vectorised techniques on rows (b) and (d) have approximately the same performance. Since the four dimensional image is set up only once, as the number of interpolations increases the technique on row (d) performs consistently better – taking approximately 25% less time in this configuration.

For comparison, an interpolation library which is part of SCIPY is also included in the comparison – row (c). This library is called `scipy.ndimage.interpolation.geometric_transform` and included in the comparison with the caveat that this library must take a single transformation matrix as the transform to apply to an image – rather than a list of coordinates. This lacks in generality, and complicates implementations. The performance of this technique is a little less than the vectorised trilinear interpolation on row (b).

	1 iteration	10 iterations	100 iterations	1000 iterations
a) Interpolation point-by-point	6.47 s	64.33 s	641.19 s	not tested
b) Vectorisation of Interpolation	0.13 s	1.13 s	11.22 s	112.62 s
c) SCIPY.ndimage (interpolation.geometric_transform)	0.13 s	1.23 s	12.45 s	127.04 s
d) Vectorisation of Interpolation with 4D image containing local neighbourhood at each point	0.13 s	0.97 s	9.00 s	90.58 s

Table 5.2: Table showing the gain in speed with the different implementations of trilinear interpolation

Table 5.2 shows that a relatively fast trilinear interpolator has been developed in this work, that takes advantage of the optimised matrix operations in PYTHON as well as the fact that a

single image needs to be interpolated many times, allowing some pre-computing to be done.

Transformation of one image

With a working 3D interpolation technique, one of the two images – the reference image in this case – can now be transformed by interpolation, in order to match it by image correlation to image in the deformed configuration.

Transformations are simply applied to an array containing the coordinates of each point in the image. This can be seen in Figure 5.5 step 5. These coordinates can be translated by any arbitrary amount (*i.e.*, not restricted to integer values) by simply adding a number to all x , y or z coordinates. Any other affine transform can be applied with relative ease to this array of coordinates, but for the purposes of this work, only rigid-body displacements and rotations have been implemented as possible grains transformations, since these are the only transformations expected, at the resolution of observation of the grains.

Rotations, as described by the axis and angle rotation vector used in this work rotate an object around the origin. Since the origin $(0, 0, 0)$ is in the corner of the array of coordinates being rotated this means that applying a rotation directly on the coordinates would cause the image of the grain to rotate around the corner of the box, and thus potentially out of the cubic subvolume. When a rotation is to be applied, the coordinates are translated in such a way that the origin is in the centre of the subvolume. The rotation is then applied, and the opposite translation applied. At this stage, any further translation desired can be added to all the coordinates.

An array of transformed coordinates can then be passed to the interpolation function developed for this work (see above in Section 5.3.6 – and step 6 in Figure 5.5). This function will, for each transformed point, interpolate a greyvalue for it based on its position in the untransformed reference image. The result is a 3D greyscale image which has been transformed with respect to the original reference image.

Image correlation

The previous two subsections (“Interpolation” and “Transformation of one image”) have described tools that allow the arbitrary, rigid body rotation and translation of a 3D image. To perform image correlation all that is needed is to systematically sweep the $x, y, z, \theta, \phi, \psi$ parameter space, generating a transformed reference image with these parameters and calculating the NCC between this new image and the untransformed, deformed image, looking for a local maximum of the NCC.

The image correlation part of this measurement has as input two identically sized cubic matrices with the two grains (in reference and deformed configurations) to correlate roughly in the middle of both matrices. This may seem like a much easier problem than that faced by TOMOWARP (continuum correlation code described in Section 5.2), where every possible combination in the displacement parameter space is explored. This is however not the case: the three additional degrees of freedom (the rotational ones) create a large parameter space to explore. The consequence of having a large parameter space is that setting a fixed step in the search of displacements and rotations involves an enormous amount of calculations if an acceptable precision (and therefore small steps in all degrees of freedom) is stipulated. As discussed in Section 5.1, the search for the highest Normalised Correlation Coefficient is typically not done with a fixed step, but rather piloted by an optimisation or minimisation routine. In this case, a minimisation routine called BFGS (as used by Bay (2008)) has been implemented. The BFGS optimiser is asked to minimise the NCC which is provided by a function called “`transformation_to_ncc`” with respect to $x, y, z, \theta, \phi, \psi$. This is shown in Figure 5.5, step 8. `transformation_to_ncc` takes those six variables as input, plus the two cubic matrices containing the images of the grains in reference and deformed configurations. The reference image is transformed by $x, y, z, \theta, \phi, \psi$

using the trilinear interpolator described above, and then the NCC between this interpolated, transformed reference image and the deformed image is calculated. Since the optimiser is trying to *minimise* a function, $(1 - NCC)$ (as defined in equation 5.1, note NCC always gives a value between 0 and 1) is returned to the BFGS algorithm.

An optimiser or minimiser like BFGS typically varies the variables which it is asked to optimise in such a way to minimise the output of the function ($(1 - NCC)$ in this case). This is typically done with gradient-based approaches that measure how much $(1 - NCC)$ increases with small changes in $x, y, z, \theta, \phi, \psi$, and then make a step in the direction of maximum gradient. It is also possible to provide an analytical gradient function to the the minimiser. This is detailed in Smith *et al.* (2002), however the formulation of the gradient of the image with respect to the rotational degrees of freedom is computationally intensive, and was not considered to add enough of a speedup for the time required to implement it.

The minimisation function continues to evaluate a gradient, and progresses in its maximum direction until a stopping criterion is reached. Stopping criteria can take many forms, but are typically expressed either as a maximum number of iterations, an error tolerance, or a minimum gradient. In this implementation a tolerance on the minimum gradient is imposed.

The output of the optimisation/minimisation procedure is the combination of $x, y, z, \theta, \phi, \psi$ which gave the minimum value of $(1 - NCC)$, as well as the actual value of $(1 - NCC)$ as well as the number of iterations performed. The typical number of interpolations is around 200 for Caicos ooid grains, and around 300 for Hostun sand grains.

Summary

The grain-based DIC approach described in this section adds a much more sophisticated measurement tool for 3D rotations of individual grains. Instead of describing a grain's orientation with a coordinate system coming from its morphology (such as from the moment of inertia, or the Feret measurements), in this technique the entire image of a grain, in two configurations is used. This implies some data processing in order to extract the grain, and the manipulation of much more data than a simple three-vector base. However, the comparison of these two techniques in Section 5.3.7 shows that the performance of grain-based DIC justifies this more sophisticated measurement technique.

5.3.7 ID-Track: experimental characterisation of measurements


To characterise the metrological performance of the different grain-based measurements developed in this work on real test cases and in conditions similar to the images acquired during the mechanical tests, some test images with known, *a priori* kinematics are required.

Rotations

A pair of images in which grains have an imposed rotation are simply acquired by imaging a specimen in a reference configuration, and then imaging it in a subsequent configuration after having applied a rigid body rotation to the entire specimen. This kind of test is performed and evaluated in Pannier *et al.* (2010). For this study a specimen is created with layers of the different types of grains studied in this work, separated by some spacer (Play-Doh[®] was used in this case). Figure 5.7 shows a vertical slice of a part of the specimen in the two configurations in which it has been imaged. The layers of each type of grain are designed in order to be several grains thick. Grains are contacting within each layer. Care is taken to be as close to the specimen size and scanning conditions used in the mechanical tests made in this work. The specimen is then inclined by a known amount and re-imaged. From the beginning, the specimen is subjected to a small amount of overburden pressure to ensure that grains do not rearrange when the specimen is tilted.

Between the two imaged configurations, each grain shares the same axis and angle of rotation as that of the specimen. Grain translations, however depend on the distance of each grain from the centre of rotation.

RTEA01 Reference Configuration



Caicos ooid

Play-Doh

Ottawa 50-70 sand

Play-Doh

Hostun HN31 sand

RTEA01 Tilted by 4.7°




Figure 5.7: Cropped vertical slices of rotation test RTEA01 in reference and inclined configurations. Note the brighter Caicos Ooids in the image; this is because these grains are calcite, compared to Ottawa and Hostun’s quartz grains

Since grains of the same layer are contacting in each 3D image, the images are treated with the procedure for identifying grains described in chapter 4, in exactly the same way that images coming from a mechanical test are treated. It is important to note that, as described in Section 4.1.2, the procedure for separating grains is not error-free which will introduce a certain amount of error in the measurement of displacements. This is particularly true of a rotation test in which the specimen rotates with respect to coordinate system of the voxels with which it is reconstructed. As Figure 4.14 shows, there is a considerable bias in the separation program used (Visilog) towards contacts in the cardinal directions of the image, consequently some changes in the segmentation are expected between configurations.

When grains are identified and measured in each state, ID-Track is then used to link grain labels in one configuration to labels in the next configuration. Once this link is made, the rotation of each grain can be measured, and statistics made on the measurements made. This allows an initial characterisation of the metrology of the measurement of rotation including the errors induced by the imaging system, as well as the subsequent steps of image analysis – in particular the separation of grains. The kind of rotation test described allows a characterisation of the total error, and not just of the two measurement techniques being compared.

Both techniques for measurement of grain rotations have been used to measure grain rotations between these two configurations. Figures 5.8 and 5.9 show distributions for all the measured rotations in this rotation test for Caicos ooids and Hostun sand respectively. The imposed rigid body rotation corresponds to a vertical line at 4.7° degrees.

Figure 5.8 blue curve, shows a wide distribution of the rotation angle measured with the moment of inertia tensor (as described in Section 5.3.5) for around 6000 Caicos ooids. The global peak of the data is actually at around 0 degrees, with a secondary peak around 5° as expected. There is a very wide distribution of measured angles which continues off the scales of the graph.

The rotations of the same grains when measured by grain-based image correlation, as described in Section 5.3.6 are shown in red. The distribution of measured rotations is much tighter

than that measured from the moment of inertia tensor – with the majority of measurement within $\pm 2^\circ$.

Figure 5.9 shows the same measurements on grains of Hostun sand. As can be seen in Figure 5.7, there are more Hostun sand grains (approximately 9000 are used in the image correlation) than Caicos ooids, which also explains the different frequencies in the histograms. The measurements coming from the moment of inertia tensor appear to be distributed approximately in the same way as they were for the Caicos ooids. The measurements coming from grain based correlation are however improved when compared to those made on the Caicos ooids, despite the fact that there are probably more separation errors when treating the Hostun grains.

Hostun grains are more angular than Caicos ooids, which are both more rounded as well as being closer to spherical. The closer objects are to spherical, the more difficult it is to measure their orientation, or their rotation between two images. Hostun grains have more eccentric character than Caicos grains, allowing the grain-based image correlation to converge more accurately.

Figure 5.10 shows the rotation *axes* measured for the Hostun sand grains analysed in this rotation test. It is clearly visible that the measured axes of rotation are aligned along the x-axis (*i.e.*, horizontally in the figure), which corresponds to the axis of rotation visible in Figure 5.7.




Figure 5.8: Histogram showing the distribution of the measured rotations for approximately 6000 Caicos Ooid grains subjected to a rigid body rotation of 4.7° . The bin width is 0.5°

As performed in Pannier *et al.* (2010), the image correlations thought to have been successful (in this case those having converged before 500 iterations and with a NCC greater than 0.985) are characterised in table 5.3. These do not compare very well to the standard deviations reported in Pannier *et al.* (2010) (between 0.12 and 0.16 for each angle of rotation), however this is certainly because – despite imaging conditions being similar – the D_{50} for the Hostun grains studied in Pannier *et al.* (2010) is approximately double than that used for these images.

Table 5.3 shows that the grain-based digital image correlation approach developed has a much higher performance in the measurement of rotations in the type of analysis used in this work.

Figure 5.11 shows the results of the application of the two methods for measuring grain rotations in an experiment where there isn't *a-priori* knowledge of the rotations expected. This said, the two configurations being compared are very close together at the beginning of the specimen's macroscopic response in which little localisation is expected. Figure 5.11 shows each grain coloured by the rotation angle from the axis and angle representation of rotation, which is considered to be always positive in this work (a negative angle is turned into a positive angle




Figure 5.9: Histogram showing the distribution of the measured rotations for approximately 9000 Hostun sand grains subjected to a rigid body rotation of 4.7° . The bin width is 0.5°




Figure 5.10: Stereoplot (see 4.2.5 for definition) showing the orientation of the axes of rotation measured for the Hostun grains in the rotation test analysed

simply by reversing the axis of rotation). The figure confirms that the measurement based on the moment of inertia tensor is very noisy: many grains have false readings of large rotations. The grain-based DIC developed to improve the measurement of rotations can be seen to do a much better job with no grain being assigned unrealistically large rotations

	Hostun Grain DIC	Hostun Moment of Inertia	Caicos Grain DIC	Caicos Moment of Inertia
Number of Grains	17703	4219	6457	2922
Mean	10.89°	4.50°	13.33	4.56°
Standard Deviation	1.16	15.71	1.51	16.64

Table 5.3: Table outlining the performance of the different two for the measurement of rotations tried in this work (using a difference of orientations calculated from the moment of inertia tensor or grain-based digital image correlation) in a rotation test where 4.7° rotation is imposed. Results are presented for two different types of sand




Figure 5.11: Slice of COEA01-01 showing grains coloured by the angle of rotation from the axis and angle representation of rotations measured on increment COEA01-01-02. On the right they are measured by the change in the maximum and minimum eigenvectors of the grain’s moment of inertia tensor. On the left rotations are measured with grain-based DIC

Displacements

In principle, the kind of test performed for the characterisation of rotations can also be used to characterised the measurement of grain displacements. It is however difficult to measure an imposed displacement with sufficient precision (compared to the expected error of less than 1.5 μm in the conditions in which the mechanical tests have been performed) in order to be able to compare the results of the displacement.

The inability to measure the imposed displacement is aggravated by the fact, as discussed in chapter 3, that the x-ray tomography setup used for this work is known to undergo mechanical fluctuations which have been related to temperature (see Section 2.3). Very briefly, this means that the effective axis of rotation is liable to move during the scan. During the reconstruction of the data this is smoothed out by choosing the centre of rotation for all images which gives the sharpest reconstructed slices (illustrated in Figure 2.8). This corresponds to a mean axis of rotation which best fits the data. The result of this, however, is that the effective centre of the reconstructed 3D images fluctuates. It has been observed to fluctuate considerably more (in the order of 10 μm) than the minimum error expected for the measurement of grain displacements.

Therefore a displacement test could generate the same histograms as for the rotations, but the actual displacement applied could not be known to a satisfactory precision. Such a test could be used only to characterise the *spread* of displacements measured, without being able to compare the value. Since a ground truth cannot be established with sufficient precision this study has not been performed.

A potential method for generating some displacement test images, where the displacement is known to sufficient precision, is to exploit the multi-scale nature of the x-ray scanner in Laboratoire 3SR, and to scan a much larger specimen, taking care however that the grains remain the same size (in pixels) in the image. Performing the scan on a much larger specimen will reduce the relative effect of the mechanical instability of the imaging setup, as well as facilitating the measurement of the imposed displacement between configurations. It will be difficult, however, to find scaled-up versions of the grains with the same morphological properties of the different types of grains tested.

Separation-free test specimens

Another kind of specimen, allowing the characterisation of kinematic measurements without incurring the penalty of potential segmentation errors, as well as the small biases possible from the tracing of the watershed lines is also possible.

This technique relies, as speculated above, on finding a satisfactory scaled up version of the grains under study. The idea is to generate a specimen of grains in 3D that *do not touch each other*. This has been achieved, for example, by the patient application of hair gel to coat multitude of large angular grains, which are then assembled in a cylindrical specimen – but which remain separated by the hair gel, whose density is low. A 3D image coming from a scan of this kind of specimen can be thresholded, in order to identify the solid phase, and in theory can be directly labelled, without a separation procedure. This is advantageous since it removes any bias from segmentation, so allows a characterisation of just the calculation of the precision of the centre of mass. This kind of specimen has not been yet tested, but has been created successfully for another series of tests on aggregates for concrete (Piotrowska, 2013 - expected).

5.3.8 ID-Track: potential improvements

Although in general ID-track has good performance when applied to the analysis of the tests that form the body of the experiments in this work, some clear areas of improvement can be identified, there are listed below with an associated “difficulty” rating.

- **Improved measurement of grain displacement (Difficulty $\frac{1}{5}$)**

If improved precision is needed in the measurement of displacements, some of the techniques for improving the measurement of binary objects described in Section 4.2.1 can be implemented. Two families of improvements are discussed (some using only the binary images, and others which bring in the greyscale information), both of which show potential for reducing the causes of poor precision in the calculation of the centre of mass. In order to choose between these techniques, the displacement tests discussed above should be accurately performed in order to be able to evaluate each technique against a known ground truth.

- **Intelligent search window (Difficulty $\frac{2}{5}$)**

An improvement to the ID-Track technique that would improve both the speed as well as the error rate, is to implement a search window that is not defined globally – as is currently implemented – but one which is defined locally. This improvement is foreseen in the current implementation of ID-Track.

A local search window is best implemented starting with some *a-priori* accurate measurements of kinematics of some parts of the specimen. These could be obtained in the ways discussed in this chapter – either by continuum DIC performed on a few nodes (as was done in Hall *et al.* (2010)), or with particle tracking – this is what the second pass of ID-Track (which penalises displacements which are far away from the mean displacement) attempts to do, however an incorrect tracking in the first step can cause this to make errors. A last option would be a manual tracking of ten or so points on the boundary of the specimen.

In any case, when a sparse set of guaranteed measurements are made, the approximate displacement field of the specimen can be interpolated at any point in the specimen. As stated in the discussion in the introduction of this chapter, a discrepancy is expected between the kinematics of material points in a continuum representing a granular medium when compared to the displacements of particles. This means that even if a field of displacements can be very accurately determined from the boundary of the specimen, a search window (*i.e.*, a tolerance on the kinematics) still needs to be implemented around this mean field of kinematics, in order to capture grains that do not, due to geometrical constraints, or localised deformation, move as dictated by the interpolated displacement field.

Furthermore, this kind of tool would facilitate a total (as opposed to incremental) analysis of a test, since the search window can grow incrementally, as the test advances – avoiding an almost specimen-wide search window when tracking from configuration 01 to the end of the test.

- **Intelligent management of “foundling grains” (Difficulty $\frac{3}{5}$)**

When stitching together grain labels from consecutively tracked increments, when a grain disappears (*i.e.*, is not tracked over an increment), but continues to be tracked over subsequent increments, it is currently added back with a new foundling grain label, as described in Section 5.3.2.

If the number of foundling grains is small with respect to the number of grains which are fully tracked it may be possible to improve the tracking of grains throughout the test in two ways. The first is to attempt to assign the original grain label to a foundling grain. This can only be done confidently if no other foundling grains are close (in space and in invariant feature) to the one under consideration. Some foundling grains are likely to be close (*i.e.*, confusable) when this procedure is being carried out – particularly in the case where two grains are correctly tracked separately until an incorrect separation in one configuration causes them to be identified as a single particle. The result of this will be that two adjacent and potentially similar foundling grains will be created in the following increment (if the separation continues to be done correctly). To help distinguish the two grains apart, the extrapolation of the grains’ kinematics over the previous increment(s) may help tell them apart. This information is more difficult to obtain and extrapolate, which is why this technique is considered difficult to implement.

The second improvement, if the first step can be done reliably, is to reconstruct the kinematics of the grain in the increment where it was not tracked by back-calculating based on its positions before and after the increment where it was not tracked. This would fill the hole in the kinematics of the specimen, which would enrich the data sets coming from this technique.

- **Physical limits on rearrangement of grains (Difficulty $\frac{5}{5}$)**

A further improvement to the local search window discussed above is to have search limits tailored for each grain based not only on kinematics, but also weighed by probability of the different geometric rearrangements possible for a given grain.

A very simple example of physical limits is that two very similar grains are unlikely to swap places, even though ID-Track may want to make assignments that appear to make

grains swap places. A scheme that allows this to be penalised may also avoid some errors. The principal difficulty in the implementation of this improvement is that there is no obvious, simple way to assign probabilities to different scenarios of granular rearrangement – which means that one rapidly gets into the domain of simulations (possibly Monte Carlo simulations).

- **Detection and mitigation of bad separation of untracked grains (Difficulty $\frac{5}{5}$)**

The main weakness of ID-Track is that it depends on a physically correct separation to be carried out on images representing each granular configuration, so that grains are individuated. As shown in the preceding chapter, the separation tool currently used offers relatively high performance but does make some errors. These are accepted, given the high-speed operation of the segmentation code used. When grains are unreliably separated in different configurations this causes them to have different invariant characteristics – which means that ID-Track does not follow them (and thankfully does not make an incorrect assignment).

The hope is that subsequent tools will improve the quality of the separation that is done between grains. If however errors are always present, this improvement is designed to help improve the separation.

The fact that a grain is not tracked can be used as a reliable indication that an inconsistent separation has been made, since in almost all other cases a grain is tracked (is a sufficiently large search window is used). When a grain in a given reference configuration is not tracked, this should leave at least one grain in the deformed configuration without an assignment. Since by far the largest number of grains that are not tracked are due to an incorrect number of markers causing incorrect separations, there is a different number of particles representing the same object in different configurations – causing ID-Track to fail to follow these grains. The most likely scenarios when failure occurs is that in one configuration there is a particle which correctly represents a grain, and in another configuration there is either a particle which corresponds to various grains fused together, or a particle which represents only a part of a real grain (see Figure 4.5 for examples of the incorrect separations). It is possible however that one configuration is over-separated and the other is under-separated, with no particle representing a grain in any of the two configurations.

The objective of this improvement is to attempt to repair the separated image, by re-performing the separation locally. The main challenge, once an incorrect segmentation is detected, is to decide whether over- or under-separation has occurred. This may be detected by a variety of different measurements. For a start the volume of the particles involved may be a good initial indication - since over-separated grains may be smaller, and under-separated ones larger than the size of the grains expected for the material. This kind of diagnosis will of course work best with a relatively small distribution of particle sizes. Other measurements may also detect this situation: for example the specific area (surface area normalised by 3D volume) of the under-separated case will be much larger than expected for an ordinary particle.

In any case, if the type of marker-placement error can be detected without errors, the space containing the particles could be extracted from the binary volume and re-processed in order to locally force the number of markers. This affects only the local area, the entire volume would not need to be re-processed. However in the case of an under-separation being repaired, care would have to be taken to use a new label for the new particle created.

This process, although very time consuming both to implement as well as to run, could potentially help to eliminate tracking errors at the source: the incorrect separation. However, since the separation process is the root of both tracking problems as well as the incorrect contact orientations measured, effort should probably go toward fixing these “upstream”

issues directly with the separation procedure. It is for this reason (as stated in the previous chapter) that work is ongoing on this front.

5.3.9 ID-Track: summary

ID-Track, as presented in this chapter is a technique for tracking grain labels from one imaged configuration to another. It depends on a reliable separation of grains in each imaged configuration, as well as relatively small increments to work optimally.

ID-Track has been extended with a discrete, grain-based digital image correlation approach, since orientations of grains calculated from each grain's shape have proven not to be sufficiently reliable for the accurate measurement of rotations. Grain-based digital image correlation is the technique that will be used in the following chapters in order to make measurements of the grain kinematics in the specimens tested in this work.

5.3.10 Mathematical tools for 3D rotations

Projection of 3D rotations onto a 2D plane

It is of interest in the study of the 3D rotations of grains involved in a shear band, to try to understand the component of rotation that the grains have in the direction of the band. For example, if the band is progressing in a well-defined direction are grains rolling in the direction of the band? Otherwise put, in a shear band whose orientation is described by a plane containing a vector pointing in the direction in which it is progressing, and an orthogonal vector pointing to its width, are the *axes* of rotation of the grains in the band co-incident with the orthogonal (width) vector?

To study the component (*i.e.*, scalar value) of rotation in a given direction, the axes and angles describing the rotations must be projected onto a plane. This projection should make the component of a grain's rotation in the plane defined by the axis of rotation as simply the scalar value of the grains rotation (*i.e.*, the component of the rotation in this plane is 1). Conversely, the projection of a grain's rotation in any plane containing the axis of rotation should give a rotation component of 0. The projection of intermediate cases should in some way reduce the angle of rotation. The scheme selected works in the following way: A 3D rotation (described in the axis and angle notation as a unit-length rotation axis, q and a rotation angle, θ) is going to be projected onto a plane whose normal is the unit vector d , in order to calculate the projected angle which will be called θ_{proj} . Vector v is obtained by calculating the "vector rejection" of q from d , *i.e.*, the component of q which is orthogonal to the direction of d .

$$v = \frac{q \cdot d}{d \cdot d} d \quad (5.9)$$

Vector v is then rotated onto v_R using the rotation described by the axis q and angle θ . This could be done by calculating the rotation matrix that corresponds to the axis and angle representation of the rotation. However for brevity, Rodrigues' rotation formula (Koks (2006)) will be used:

$$v_R = v \cos(\theta) + (q \times v) \sin(\theta) + q(q \cdot v)(1 - \cos(\theta)) \quad (5.10)$$

v_R is then projected back onto the plane, defined by d , giving v_P . This is done, as above by calculating the vector rejection of v_R from d

$$v_P = \frac{v_R \cdot d}{d \cdot d} d \quad (5.11)$$

Finally, the projection of original angle θ is calculated by measuring the angle between vectors v and v_P

$$\theta_{proj} = \arccos\left(\frac{v \cdot v_P}{|v||v_P|}\right) \quad (5.12)$$

As stated above, when the axis of rotation and the normal to the plane of projection are co-incident, $\theta = \theta_{proj}$. This particular case needs to be detected, as it is not handled by this formulation. As the angle between d and q increases from 0° ($d = q$) to 90° ($d \perp q$), the ratio θ_{proj}/θ (thought of as a multiplicative reduction factor of θ) goes progressively from 1 to 0, as desired. Figure 5.12 shows how this ratio evolves.

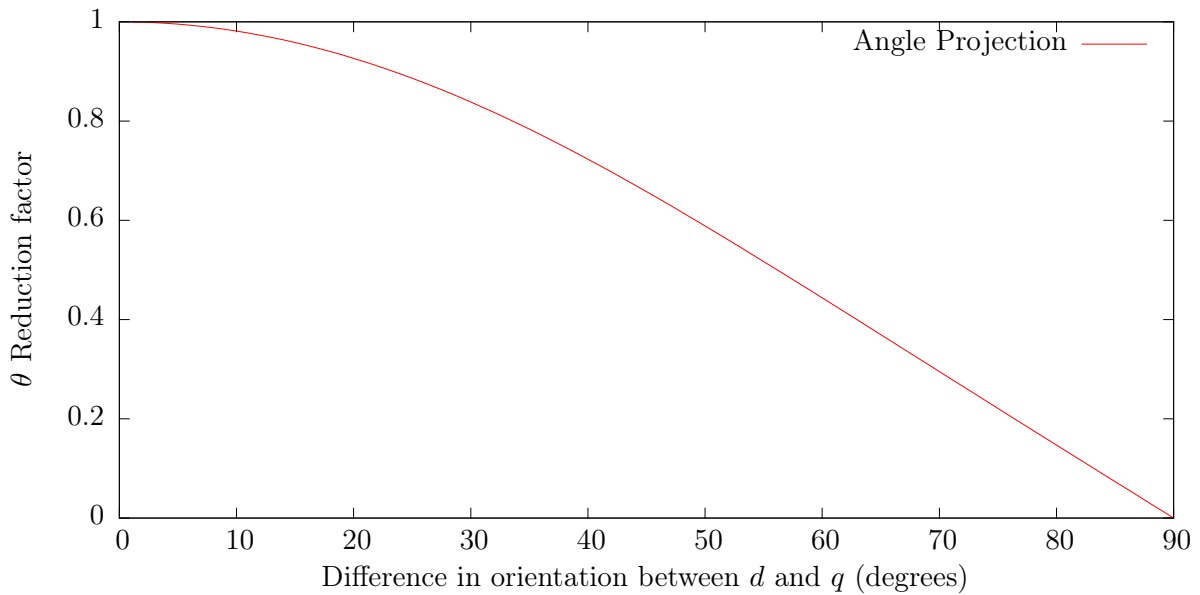


Figure 5.12: Graph showing, as the angle between a rotation axis (q) and the normal of the plane onto which it is being projected (d) increases from 0 to 90° , the reduction in the factor that multiplies the angle of rotation θ to give the projected angle θ_{proj}

This formulation is used for projecting axes and angles of rotation, when they need to be studied in a plane.

Co-rotations

Given that grain rotations are observed, an interesting mechanical process to measure could be grain co-rotation: the amount that one grain causes neighbouring grains to rotate. In the two dimensional case of grain rotations, these kinds of measurements can be made by simply comparing rotation angle scalars, which fully describe the rotation in 2D. Typically clockwise or anticlockwise rotations are written as positive and negative values of rotation. This allows one to see whether the clockwise rotation of one grain is transmitted to another grain as an anticlockwise rotation in a cog-like system.

In order to extend the co-rotation concept to 3D components of rotations must be projected onto a 2D plane. The tools described immediately above (Section 5.3.10) give a method for projecting 3D rotations onto a plane. A typical plane of interest for looking at rotations in a shear band could be the plane containing the axis of the specimen and the normal to the plane containing the shear band. When compared simply to the scalar angles of rotation this would allow one to answer questions such as: “Are grains principally rotating in the direction of the

shear band?” Otherwise, for very local analysis of rotations it could be conceivable to extract tight rings of an even number of contacting grains with similar values of rotation, and project their rotations on the plane containing them all. This could show whether this ring of grains are rotating together like a system of gears.

5.4 Contact kinematics

Once grain kinematics can be established successfully, measurements might be made of the evolution of the contacts between grains. Contacts are the points at which the forces are passed through the granular system; the orientation of contacts is inextricably linked to mechanisms of granular rearrangement such as force chain buckling.

5.4.1 Algorithm for calculation of contact kinematics

To best study contact evolutions, the distribution of contacts should be analysed both at the specimen level, and locally by following contacts individually in order to identify important events in the deformation of the specimen. Unfortunately, when contacts are defined as the watershed lines used to separate grains in the 3D images, the orientations resulting from this process are flawed: they are strongly affected by artefacts (Figure 4.14 showed the artefacts resulting from two different watershed algorithms used).

Since contact orientations cannot yet be measured satisfactorily, preliminary results showing evolution of contact orientations are of limited interest. However, as described in Section 4.2.5, a new separation technique is currently being tested. Once this is in place, contacts will need to be followed through various configurations, just like grains were in Section 5.3.2. There is, however, a fundamental difference (as far as tracking is concerned) between grains and contacts: grains are permanent objects in these tests, whereas contacts are ephemeral and can be created as well as vanish. The objective is thus to label each contact in a way that is consistent through the different configurations in which the material has been imaged. The desired result of this procedure is a large array of contact labels, that records in which configuration each contact appears, as well as a lookup table (as was done for the grains), that allows a given contact label in any configuration to be looked up to give that contact’s label in the large “contact tracking” array. To achieve this, the algorithm developed for the following of grains must therefore be modified to account for the fact that contacts by their nature appear and disappear (unlike grains which are not meant to disappear, and are called foundling grains when re-appear in order to preserve some continuity in the data). The result is a hybrid of the grain following algorithm with the algorithm for labelling contacts in a single 3D image, described in Section 4.1.4 (and Figure 4.8).

The inputs to this procedure are: 1) the output of the grain-labelling procedure, for each configuration (*i.e.*, for each grain in each configuration, the grains that are in contact with said grain, and the label given to this contact) and 2) the lookup table resulting from the whole-test tracking (which allows a grain label in any configuration to be looked up to the corresponding grain label in the initial configuration). The “contact tracking” matrix that will be output at the end has “number of contact labels” rows and “number of configuration” columns - the objective being to check off in this matrix the configurations in which a contact has been found. The algorithm bases itself on the labels assigned to contacts and grains in the initial configuration.

The “list of lists” structure described for the labelling of contacts is reproduced in memory for the initial configuration, from the data saved from the contact labelling procedure. For subsequent configurations, the same structure is aimed for, but cannot simply be directly reproduced. Rather, the data describing the contact labels in each subsequent configuration ($1 + N$) is read into memory, and each element is processed as shown in Figure 5.13: for each grain, each contact is read out, the two grain labels involved (A and B) are copied, and the contact label L is checked

against a “done” list to avoid repetitions. A and B are then looked up to their labels in the initial configuration (A' and B'). The labelled contacts list in the initial configuration is then searched in order to see whether a label has already been assigned to the contact between these two grains. If a label exists for this pair of contacting grains, it is called L' , and a 1 is put in the “contact tracking matrix” at position $L' : 1 + N$. Otherwise a new contact label (L'_{new}) is created, and is added to the labelled contacts list in the initial configuration as well as causing a new row, and a 1 to be added to the “contact tracking matrix” in position $L'_{new} : 1 + N$. At the end of the procedure the “contact tracking matrix” is written to file.




Figure 5.13: Flow chart showing a part of the label tracking algorithm

The “contact tracking matrix” is then post-processed from a matrix of boolean values into a “contacts lifetime matrix” with a simple further algorithm which for each contact, in each configuration, calculates two numbers: if the contact is currently active, how long it has been active for (if inactive this value is 0); if the contact is currently inactive, how long it has been inactive for (if active this value is 0). These pairs of values are saved for each contact, as a contacts lifetime matrix.

5.4.2 Analysis of contact kinematics

The “contacts lifetime matrix” coming from the algorithm described immediately above, describes the apparition and disappearance of contacts for each test. It allows global trends – such as longevity of contacts, and the determination of macroscopic events that cause a contacts to be broken or formed – to be characterised. Furthermore, local analysis (both in space and time) of individual contacts is also possible. Local analysis of contacts is of particular interest for making experimental observations around force chains (keeping in mind that intergranular force information is not available in these tests). An objective, however is to characterise contact behaviour and relate this to observations (both from simulations as well as experiments with local force measurement – most notably photoelastic disks with 2D loading) of force-chain buckling events. This is a subject which will be developed much further in the future, pending on higher-quality contact orientation measurements, but the tools developed will be presented for completeness.

Contacts between grains are expected to be able to carry a certain amount of “shear” or “tangential” force in the plane of the contact (as a function of the friction between grains, and the normal force in the direction of the contact). The kinematics that a single contact (defined simply as any point of contact between two grains) can undergo are divided into two categories. The first is rigid body kinematics of the two-grains-and-contact system which cause the contact to displace and change orientation in space. The second category of contact kinematics describe (shown in Figure 5.14) what happens to a contact when there is relative movement of the grains in contact. Contacts can a) slide in the direction of the plane of the contact (when the tangential force is sufficient to overcome friction) and b) rotate when there is relative rotation of the two grains, while remaining in contact (like two cogs). Twisting of the contact is also possible (Figure 5.14 c), when grains rotated around the normal to the contact. This kind of information is derived relatively easily by looking at contacts which are persistent over a given increment, and using grain kinematics calculated over the increment in order to calculate these contact kinematics.

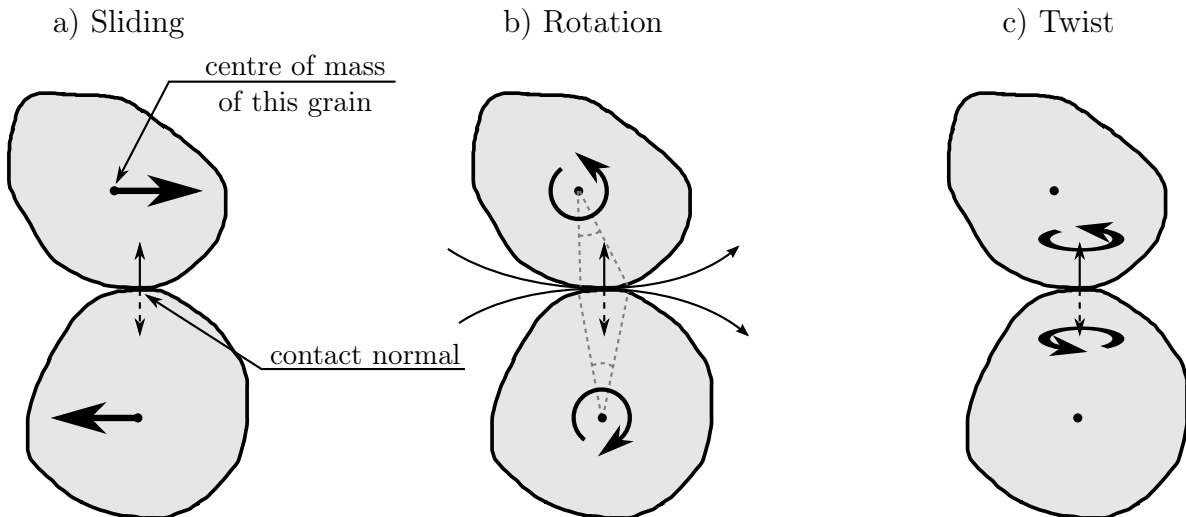


Figure 5.14: Illustration of the three kind of contact kinematics expected. Part a) describes relative sliding at the contact, b) shows cog-like co-rotation of the particles, and c) shows “twisting”

The calculation of the contact kinematics resulting from differential grain kinematics shown in Figure 5.14 are relatively easily obtained for situations a) and b) by calculating the kinematics of each grain at the point of contact (*i.e.*, adding the displacement vector due to the 3D rotation at the middle of the contact to the rigid-body displacement of the grain). Twisting (situation c) is calculated by projecting the 3D rotation of each grain onto the contact normal by using the tool shown in Section 5.3.10.

Since the direction of the major principal stress is well-defined in these tests an interesting

measurement may be to find networks of grains who contact each other at an angle within a certain tolerance of the direction of principal stress, and too look at the longevity of these structures.

Another particularly interesting and more local prospective is to find contacts which are persistent during the beginning of the test which are then lost, to relate this loss-of-contact to event to the different types of contact kinematics described above.

5.4.3 Conclusion

Although the tools are in place, this analysis is currently flawed due to the observed errors in the measurements of contacts between grains. Although the presence of contacts between grains can be assessed with confidence at around the spatial resolution of observation the orientation of the contacts coming from the separation lines of the watershed algorithm are too biased to use for significant measurements. In the near future further work will permit the accurate measurement of contact orientations, allowing these tools to be used on the data collected in this work.

5.5 Micro-strain from grain kinematics

This section briefly introduces local measurements of strain which can be made with grain kinematics. This subject is the object of two comprehensive papers: Bagi (1996) “Stress and strain in granular assemblies” formalises a convenient description of stress and strain at the particle level – several preceding publications, summarised in Satake (1989) also make contributions in this domain. Thereafter, Durán *et al.* (2010) make a review of different definitions of micro-level strain definitions comparing Bagi’s “equivalent continuum strain” definition to three other so-called “best-fit strain formulations”; particle-, contact- and edge-based. The strain definition from Bagi is considered as one of the more complex formulations, however it is found to perform well.

All the different strain calculation techniques try to calculate a local strain tensor based on the neighbourhood displacements of surrounding particles. Bagi’s formulation requires a triangulation of the grain centres, to identify neighbouring grains.

In this work these definitions have not been re-implemented since an interface module called “yade.wrapper.TesselationWrapper” was graciously added by Bruno Chareyre to the open source finite-element code Yade (Šmilauer *et al.*, 2010). This interface module allows particle kinematics coming from ID-Track to be input into the in-built micro-strain calculation in Yade. The details of the micro-level strain are available in Catalano *et al.* (2013).

3D micro-strain volumes are created in this way, allowing a local measurement of strain to be made for some key increments analysed in this work.

5.6 Conclusions

This chapter has introduced digital image correlation as a procedure for making measurements between two images of deforming granular media. The chapter briefly described a continuum image correlation code (TOMOWARP) developed by Stephen Hall, which has been used for some analysis in this work. The main part of the chapter has been dedicated to grain-based measurement techniques. A previously developed technique used in Hall *et al.* (2010) was described and introduced; however the focus of the chapter is the technique developed in this work: ID-Track. This technique requires two separated and labelled 3D images in order to work, and its objective is to make a link between labelled grains in each configuration. Once grains are linked between configurations, the displacement of each grain is easily calculated by following its centre of mass. Grain rotations are more difficult to measure accurately, since the standard methods for orienting particles in a given configuration are not sufficiently reliable for this application. Inspiration was

therefore taken from Hall *et al.* (2010), and a discrete (grain-based) digital image correlation scheme was developed from the ground up.

The performance of the measurement of grain rotations was evaluated on a specimen of all different types of grains, which was imaged in two different configurations between which the specimen was rigidly rotated. This comparison shows substantial improvement in the measurement of rotations using the grain-based image correlation approach compared to the difference in grain orientation. Furthermore, even when considered net of all the possible upstream errors, the rotations measured appear to be accurate to the nearest degree or so, meaning that this tool can be used to make precise micro-mechanical measurements, which will be presented in the following chapter.

The tools required for the analysis of contact kinematics have also been described in this chapter; although results are not presented, since they depend on accurate contact definition which is currently unavailable.

Finally table 5.4 shows the approximate computation times for a typical 3D image analysed in this work.

	Calculation Time
Grain tracking with ID-Track	10 minutes
Calculation of micro-strain with Yade	1 minute
ID-Track grain-based image correlation	4 days
TOMOWARP tracking of a single increment	f(node spacing) Many tens of hours

Table 5.4: Table showing approximate processing times for a single greyscale image coming from tomography

Chapter 6

Results and discussion

This chapter presents examples of the micro-scale measurements made possible with the tools presented in Chapters 4 and 5. This chapter will start by focussing on a single test, COEA01, a triaxial compression test on Caicos ooids performed at 100 kPa confinement pressure. This will then be followed by a discussion of this test with relation to the others, in order to quantify the difference that grain type and confining pressure have on the micro-mechanical behaviour of granular materials tested.

All the analysis performed in this work is in 3D, which means that communicating these data, in 2D print is challenging. Vertical cross sections (also called slices) through scalar fields will principally be shown. The scalar values of the field in question will be coloured by a colour map, which will accompany each image. Unless stated otherwise, the vertical slices shown by default will be oriented to include both the axis of the specimen, and the normal to the shear band that develops by the end of the test. In this orientation the shear band will cross the slice diagonally from one side of the slice to the other.

The objectives of this chapter are not to make a single specific measurement, but rather to use a panoply of tools to characterise the micro-mechanisms imaged by tomography at a variety of different scales. More specifically, the objective is to use these tools in the characterisation of the onset and development of shear banding in granular materials.

Layout of chapter

This chapter begins by briefly recalling the macroscopic behaviour of the specimen COEA01. Section 6.2 then shows some of the “state-wise” measurements detailed and developed in Chapter 4 for each state that the specimen has been imaged in. Several types of state-wise measurements are highlighted, with a progression towards increasingly local measurements; where the high-resolution images acquired in this work really allow some new observations to be made.

The first type of measurement discussed is at the specimen-scale. Measurements of the bulk volume of specimens are inherent in the volumetric strain plots shown systematically in the Chapter 3. The local, mesoscopic measurement of porosity (calculated on a local REV) is then illustrated. This is followed by measurements at the grain scale, which are used to reconstruct a grain-size distribution for different steps of a test to confirm that little or no grain breakage is happening. Grain-scale measurements also allow the orientations of grains to be characterised. At an even smaller scale, local measurements of contacts (more specifically coordination number) will be made, and analysed throughout test COEA01.

The chapter continues onto the kinematic analysis with the tools developed in Chapter 5, showing grain displacements and rotations in increments throughout the test. The derivation of a strain field allows non-homogeneous displacements to be identified more easily. Several techniques exist for deriving a strain fields from discrete grain displacements, the one used in this work is briefly described in Section 5.5. The strain fields obtained using this technique are

compared to the strain fields obtained by Continuum DIC (using TOMOWARP). The principal observations made in on specimen COEA01 are then summarised in Sections 6.4 and 6.5 of this Chapter.

After the analysis of a single test, Section 6.6 analyses the other tests performed in the experimental campaign, comparing them to the observations made on COEA01. The objective of this part of the chapter is to see whether different micro-mechanisms can be identified in tests with different grain types and at different confining pressures. Once different micro-mechanisms can be identified, the challenge thereafter becomes the explanation of the differences macroscopic behaviour of the specimens using the observed micro-mechanisms. Section 6.6 is divided into four different parts: Section 6.6.2 attempts to explain the difference between two backup tests, which have considerably different peak stresses. Thereafter Section 6.6.3 looks at the micro-mechanics of the fully-developed shear bands in the different specimens tested in this work. Based on those observations, Section 6.6.4 then looks earlier in the tests in order to see the micro-scale development which lead to shear banding. Section 6.6.5 focuses instead at the very first increment of each test, in order to evaluate the initial homogeneity of the micro-scale responses of each specimen.

6.1 Experiment COEA01 ($\sigma_3 = 100$ kPa)

The macroscopic measurements made on specimen COEA01 is briefly recalled in this section, although they are discussed in detail along with all the other experiments in Section 3.6. Figure 6.1 shows two vertical slices taken from the 3D image of the specimen in state COEA01-01; these slices are taken are 90° from each other and the slice on the left is oriented to show the maximum tilt of the bottom platen from the horizontal (measured as 0.6°). These slices show that the specimen is essentially vertical – almost no deviation from verticality can be measured from the boundary of the specimen.




Figure 6.1: Slices of specimen COEA01 in state 01 (*i.e.*, under isotropic compression of 100 kPa), enlarged versions of slices shown in Table 3.5

Figure 6.2 shows the typical stress-strain responses measured for this test (as already shown in 3.15). The axial stress response, normalised by the confining pressure (which is constant in this case) is plotted against axial shortening (top). The different stages at which axial loading is

interrupted to perform an x-ray scan can clearly be seen as axial stress relaxations. The different granular configurations that have been imaged are numbered on this plot. The volumetric strain is also plotted against axial shortening (bottom). The volumetric strain is calculated from the 3D images acquired. The specimen has a qualitatively typical response for a dense specimen: there is a clear peak in the axial stress response, followed by strain softening into a residual stress plateau. The peak in the axial stress response is at a stress ratio of 6.79 (this corresponds to a high friction angle of 48.0°), which occurs at a shortening of 3.66%. A residual stress ratio of 3.70 (friction angle 20.3°) is achieved starting from 11% shortening, although by 8.5% axial shortening the axial stress is already close to the value achieved in the final plateau. The specimen's volumetric strain response is consistent with what is expected for a dense specimen: there is a small contraction between configurations COEA01-01 and 02; thereafter the specimen dilates until it reaches its residual stress state. Over the increment COEA01-04-08, which contains the peak, the angle of dilatancy is 20° .

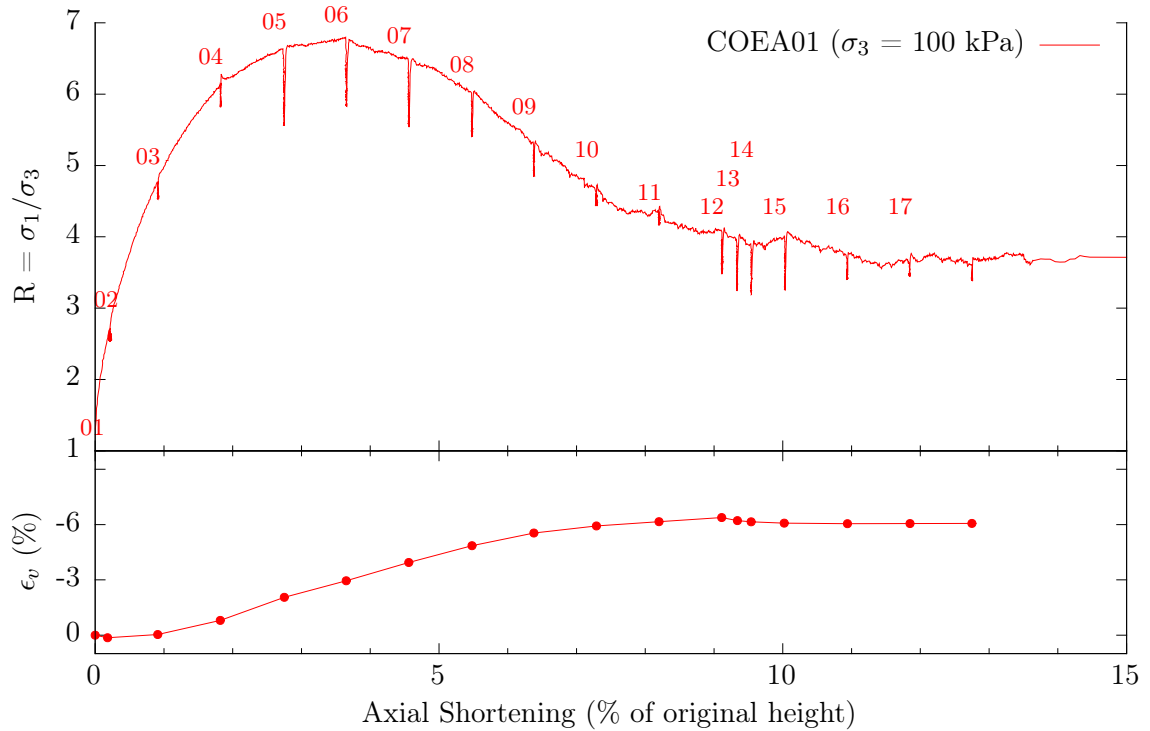


Figure 6.2: Macroscopic response of specimen COEA01

6.2 COEA01: analysis of single configurations

This section shows some key measurements possible on single images (*i.e.*, without tracking) of the test COEA01. The objective is to show the kinds of measurements that can be made on the 3D images coming from tomography, and where possible comparing them to measurements obtained by other means. The section starts by showing specimen-level measurements, and progressively zooms in, through REV-based measurements (porosity) to grain and contact-based measurements (both of which require images to be separated and labelled).

6.2.1 Global measurements: specimen-scale

The specimen-scale measurements possible in the 3D images coming from tomography are not very numerous. The principal kind of measurement that can be made is the characterisation of the (external) geometry of each specimen.

Chapter 3, which details the experiments performed in this work, systematically shows some simple measurements of the each specimen’s geometry in its initial state (its length, diameter and angles of tilt of the specimen’s boundaries). The measurements of each specimen’s initial condition have been made on the 3D images coming tomography. Furthermore, the volumetric strain for each specimen has been calculated in all its imaged configurations by computing its bulk volume using the technique shown in Section 4.4.

6.2.2 Mesoscopic measurements: local porosity

The volumetric strain (calculated from the specimen-scale measurements of bulk volume) against axial stress graphs shown in Section 3.6 clearly show an increase of the specimen volume during these tests, which means that porosity must also be increasing (since grain are not expected to undergo much volumetric strain in these conditions).

Figure 6.2 shows the macroscopic specimen response for specimen COEA01. The configurations that have been analysed are indicated on the top part of this figure. Figure 6.3 shows slices taken through the calculated porosity fields, where each pixel represents the value of porosity calculated on an REV measuring 28^3 voxels – a value approximately equal to the dimensions of a grain, calibrated in previous work to give the best results on these images. The spacing between calculation nodes is 10 pixels, so there is considerable overlap of measured volumes. The 3D porosity image shown have been filtered for noise with a 3D median filter of radius 3 pixels (which simply replaces each voxel’s value with that of the median value of voxels inside a very small sphere centred on that voxel).




Figure 6.3: Figure showing slices through the measured 3D field of porosity in different configuration of the test COEA01

The porosity slice shown for the initial state (COEA01-01) shows that the porosity of the specimen is relatively uniform, with the majority of readings in the range of 32.5 to 37% porosity, with a peak frequency of 34.9%. This corresponds well to the overall specimen porosity reported in table 3.3 of 35.2% The porous stone is not visible in these images because compared to the relatively dense grains, it is not selected by the threshold applied to the greyscale images.

For the purposes of illustration, and also in order to underline the fact that all the measurements presented herein are 3D in nature, the 3D field of porosity measured for the last increment in Figure 6.3 (COEA01-17) is rendered as a 3D image in Figure 6.4. The convenient way in which measurement nodes are regularly spaced means that the resulting fields can be interpreted themselves as 3D images is explained in Section 4.3.1, and in particular in Figure 4.21.




Figure 6.4: 3D rendering using Paraview of the porosity field of specimen COEA01 in configuration 17. A vertical portion has been cut away to reveal the inside of the specimen

Referring back to Figure 6.3, almost no changes in the distribution of porosity can be detected between configurations COEA01-01 and COEA01-03. Between these two states, the specimen has slightly shortened, however these configurations are both sides of the contractive stage of the specimen’s volumetric deformation, and there is consequently almost no volumetric strain between these configurations. Thereafter until COEA01-05 few changes in the distribution of porosity are visible. In configuration COEA01-05, some increase of porosity is visible on the left of the specimen, following this back to COEA01-04 initial indications of it can be seen. Starting from configuration COEA01-06 (a configuration practically at the peak in the axial stress response of the specimen) indications of a band of increased porosity crossing the specimen from bottom left to top right can be noticed, with porosity increased more significantly on the left and centre of the slice.

This higher porosity band develops with porosity increasing over the subsequent configurations until COEA01-12 which correspond to the residual stress state, at this point the band appears to have reached a stable value of porosity. The thickness of the band does not appear to change as it develops. This band is consistent with a dilatant shear band expected in a dense granular medium being sheared in these conditions (particularly under relatively low confinement).

From configurations COEA01-08 onwards, two distinct regions can be identified within the band, with the top having a higher porosity than the lower. As will be confirmed later in the

chapter with the kinematical measurements, the band with the highest porosity remains active and can be seen to increase in porosity throughout shearing whereas the lower part of the band does not.

The development of porosity in COEA01 measured along a profile perpendicular to the shear band as it evolves is shown in Figure 6.5 (note that COEA01-01 is in blue, and subsequent profiles become increasingly red as the test develops). The profile goes from the top left of the specimen to the bottom right. It is worth underlining that since the top of the specimen has zero imposed displacements, the porosity in the profile towards the top of the specimen coincide much more closely than those at the end. This figure goes some way to prove that the band of increased porosity seen in the slices does not appear to concentrate (*i.e.*, become thinner as it progressively develops), instead a relatively well-defined region of the specimen has an increase of porosity as the test develops. A maximum value of around 49% porosity is achieved on a few nodes by the end of the test and, as can be seen in the slices, these values can be found all the way through the volume inside the band.

The band of increased porosity starts from around 33 pixels along the profile and ends at the end of the profile at 90 pixels, with a region of particularly high values of porosity between 41 and 70 pixels along the profile. This band of particularly high porosity is therefore spanned by 29 nodes which corresponds to 290 pixels and thus, at the pixel size of this image $29 \times 15.556 = 4511 \mu\text{m}$ which is approximately $10.7 \times D_{50}$ between the centres of the calculation nodes.

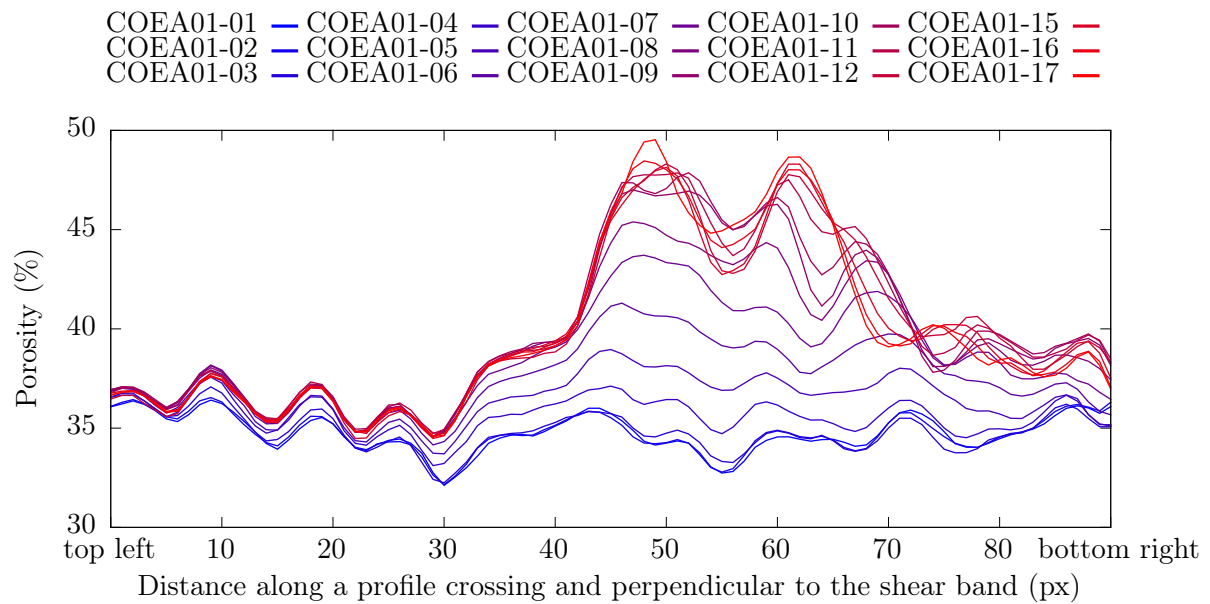


Figure 6.5: Evolution of measured porosity in specimen COEA01 along a profile which is perpendicular to the band of increased porosity in the final state. Note that line colours go from blue to red as the test develops. Configurations COEA01-13 and COEA01-14 have been omitted from this plot since they represent very small changes from COEA01-12

The nodes involved in this band can be identified with relative ease in 3D and once identified, their porosity measurements can be averaged in order to give a mean porosity for the band. The 3D zone inside the specimen that the dilatant band occupies by the end of the test can also be studied in earlier increments. Figure 6.6 (blue line) shows the evolution of the mean porosity in the region of the specimen occupied by the shear band specimen by the end of the test.

As expected, Figure 6.6 shows both the porosity in the shear band as well as the specimen's overall porosity tending towards different final values. The value of porosity in the shear band is considerably higher the specimen's overall porosity, which confirms the distributions of porosity shown at the beginning of this section. The trends in average porosity are not surprising; the shear band is a relatively large phenomenon compared to the size of the specimen. By the end

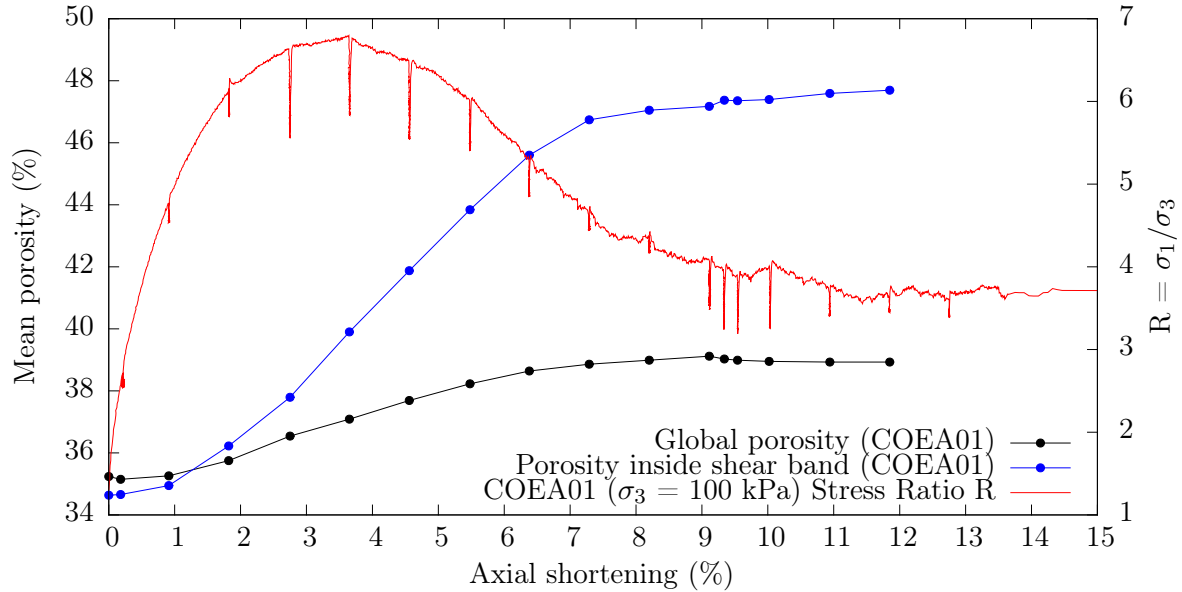


Figure 6.6: Plot of the mean porosity of both the specimen COEA01 (black line) and exclusively in the zone of the shear band (blue line) against axial shortening. The specimen’s stress ratio is given in red, also against axial shortening

of the test the shear band is expected to be the only part of the specimen undergoing strain, which means that as it increases in porosity, the overall porosity of the specimen increases.

These results are reminiscent of those presented in Desrues *et al.* (1996), where global and local void ratios in Hostun sand are studied by x-ray tomography, showing that void ratios in shear bands tend very closely to the same final values. A similar study is performed on the much smaller Hostun sand specimens studied in this work later in this chapter.

6.2.3 Evaluation of initial inhomogeneities

Intuitively, the reasons behind the material localising in a certain zone can probably be found (at least with the densely packed, and relatively small number of grains tested) in inhomogeneities in the specimen. These are expected to locally concentrate forces, causing strain. Thanks to the dense packing in these specimens, local strain will cause a local increase in porosity, which is expected to cause a cascade of further inhomogeneities that lead to strain localisation and subsequently failure.

One might expect that the such inhomogeneities should be looked for at a scale larger than the grain scale, meaning that the measurement of porosity is a good place to start. Porosity is by definition a meso- or macroscopic measurement (measuring the inside a solid grain, or inside a pore is relatively meaningless). This mesoscopic measurement is particularly well-suited to the evaluation of specimen inhomogeneities at a scale larger than the grain scale.

Without the grain-scale kinematics that are shown in Section 6.3, only the results of inhomogeneities can be searched for in the porosity fields presented in this section.

Searching in 3D for existing weaknesses to verify this hypothesis, does reveal some higher porosity areas, on the boundary of the specimen, which are connected by the shear band as it develops. Figure 6.7 highlights the potential nucleation sites of the shear band for this specimen. These zones can only be identified with the benefit of hindsight since similar zones exist elsewhere in the specimen, which do not attract the shear band.

Systematic analysis of other specimens shows that, even with the benefit of hindsight, it is not always possible to identify zones of higher porosity in a specimen’s initial state that are spatially correlated with the position the shear band that develops by the end of the test.

COEA01-01




Figure 6.7: Slice from the 3D field of porosity of specimen COEA01 in its initial state (COEA01-01). A zone of increased porosity in this configuration coincident with the zone in which the shear band develops by the end of the test is circled in black. The slice is taken in a different position than the one presented in Figure 6.3. Note, furthermore that the lookup table used to colour the field is tighter, in order to emphasise differences in the distribution of porosity in this initial state

It may be possible to look for spatial correlations of porosity in a specimen's initial state with the position of the shear band in the specimen's final state, using an objective measurement of spatial correlation; this has not been investigated for the moment. However, the predictive ability of the porosity field of specimen's initial states appears to be limited by itself.

Conclusions

This section has shown the sorts of analysis which is possible with 3D field measurements of porosity. The 3D field of porosity calculated in the initial state of specimen COEA01 allows a 3D assessment of the distribution of porosity in the specimen before shearing. Furthermore, this 3D field can be compared to and checked against the specimen's global porosity – the comparison is good in this case.

The porosity field has also been studied during deviatoric loading, showing that the shear band that develops in the specimen also corresponds to a zone of porosity reduction. This is consistent with what is expected in a dense sand specimen in which strain localises.

The evolution of global porosity and porosity inside the shear band zone, in the style of Desrues *et al.* (1996), qualitatively confirming the type of response for a dense granular arrangement.

Some indications that inhomogeneities in the distribution of porosity are correlated in space with the final position of the shear band have been detected, however they cannot be detected in all tests, and are far from being unique.

6.2.4 Micro-scale measurements: grain sizes

When all the individual 3D images acquired during a test can be segmented, measurements of grain properties can be followed during a test. For example, grain sizes can be measured to confirm the grain-size distributions obtained by other techniques, or given by the manufacturer about the three materials tested. Unfortunately, the measurement resulting from a classical measurement of a grain-size distribution is not very clear from the point of view of grain morphology. The shaking of three dimensional grains through an almost two-dimensional mesh means that grains pass through or are retained on a given mesh size based on some features of their cross section.

It may be possible to construct measurements that reproduce the shape-based filtering that occurs when measuring a grain-size distribution in reality. These measurements would have to generalise the Feret measurements (measurements of a 3D particle's maximum and minimum diameters – Feret measurements are briefly described in Section 4.2.4) into a two-dimensional measurement.

For the purposes of this demonstration, the equivalent sphere diameter of the grains measured in the 3D images will be used as a grain-based measurement of grain size; this measurement is simply the diameter of the sphere with the same 3D volume as the grain – this therefore requires a separated and labelled volume. The distribution of the equivalent sphere diameter will be calculated in an increment at the beginning of the test, and compared both to the grain-size distribution for this material, as well as to the equivalent sphere diameter for the last image acquired.

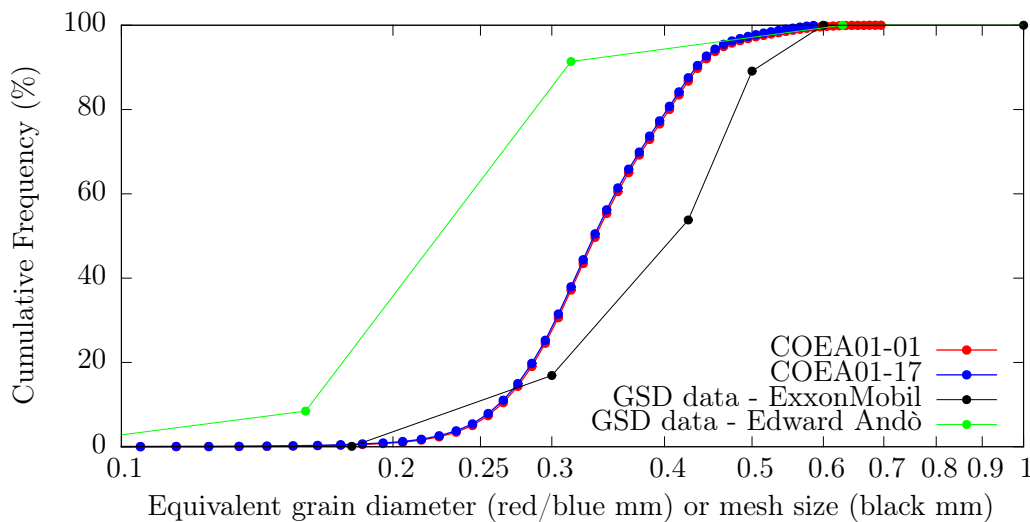


Figure 6.8: Figure comparing particle sizes for Caicos ooids, calculated from Grain Size Distribution from sieving (black – data provided by ExxonMobil, and green as reproduced in Laboratoire 3SR) and equivalent sphere diameters calculated from configurations COEA01-01 and COEA01-17

The red and blue curves shown in Figure 6.8 show the distribution of the equivalent sphere diameters calculated from two labelled images of the specimen COEA01 in two states COEA01-01 and COEA01-17 representing the end of the isotropic compression and the end of the shearing (at 11.85% axial shortening) respectively. The comparison of these two plots shows that there are negligible differences in the grain-size distribution (the grain-size distribution in the last increment appears to describe slightly larger grains, but this is negligible). This confirms the assumption that grain breakage is not happening on a significant scale in the tests performed in this work.

The comparison of the grain-size distribution calculated using sieves (black and green lines) to the data coming from tomography (blue and red) is moderately encouraging. It is immediately

noticeable however that the grain-size distribution as measured by a sieve analysis has a resolution dictated and limited by the set of sieves used whereas the data from tomography is simply a histogram of the volumes of all the grains coming from the grain separation and labelling phase. Overall the trends compare quite well, however the D_{50} of the data coming from the sieving method is larger than the D_{50} measured from tomography. It is not known how the sieve analysis was performed in ExxonMobil, therefore potential errors are difficult to pinpoint. The data coming from tomography is principally vulnerable to incorrect thresholding, which has been limited in this case by making a measurement of the mass (and therefore volume, by using an assumed grain density) of the grains making up the specimen. The other source of error is the grain separation algorithm used. Although errors are possible they are particularly rare for this rounded material, so the trends averaged over the specimen are expected to be relatively consistent.

6.2.5 Micro-scale measurements: grain orientations

Another key measurement that can be made on individual grains is the measurement of their orientation. Section 4.2.2 outlines two principal methods with which a 3D *particle* can be oriented. Some caveats with respect to the reliability of the orientations that are measured in this way are given, however they are mostly problematic when they are specifically compared for the same grain in two different configurations. The measurement obtained on many grains are expected to show the statistically correct trends, *i.e.*, the errors are expected to average out.

Figure 6.9 shows equal-area stereoplots of the orientations of the major principal axis of a set of grains whose centres lie within a 3.1 mm vertical band through which the shear band crosses, containing roughly 8000 grains. Orientations are shown for in the first and last configurations of the test. The point which defines the orientation of the shear band in the case of specimen COEA01 is in the radial position 45° left and angular position of 0° .




Figure 6.9: Orientation of grain major principal axis in two configurations of specimen COEA01, taken from a subselection of the specimen where the shear band crosses. Data are “binned” into small sectors. The plots show the number of points per bin, normalised by median number of points per bin

The comparison of the stereoplots for each material does not reveal clear preferential alignment in either state, although there appears to be a relative loss of grains with orientations

in the vertical direction (in the direction of the major principal stress), however no change in orientations in direction of the shear band.

This will be corroborated by evidence shown for grain rotations in Section 6.3, which shows considerably more disorganisation in grain rotations is measured than what is expected from 2D simulations, and experimental results from $1\gamma 2\varepsilon$ on 2D materials.

6.2.6 Micro-scale measurements: evolution of contacts

Another key measurement that can be made on a single 3D image coming from tomography (after binarisation and separation – see Chapter 4), is to characterise the contacts between grains.

As stated in Section 4.2.5, the orientation of each contact cannot yet be satisfactorily measured in the images acquired in this work (*i.e.*, with the small contact areas shown in 4.15). A project to improve these measurement is ongoing which shows promising results. The presence of contacts can however be established with relative confidence, and this can be followed through each configuration of a given test. One of the natural outputs of the algorithm presented in Section 4.1.4 for labelling contacts in a separated 3D image, is the coordination number of each grain (*i.e.*, number of other grains that each grain is touching). Figure 6.10 shows slices through selected configurations of specimen COEA01 where each grain is coloured by its coordination number. The specimen configurations as Figure 6.2 are presented. In the slices shown, the grains on the boundaries of the specimen can be seen to have a lower coordination number, as expected. Furthermore, some particularly large grains have some high coordination numbers. As the test progresses, contacts are progressively lost, although until COEA01-04 this is difficult to spot in the sequential images shown. In configuration COEA01-04 on the left, middle of the slice a zone containing grains that have a reduced coordination with respect to previous configurations can be seen. This zone extends across the specimen in the subsequent configurations and corresponds to the shear band. The coordination number reduces in the band until configuration COEA01-10, at which point the zone of reduced coordination number seems stable. The decreasing coordination number is consistent with the observed increase in porosity.

Figure 6.11 shows a histogram of coordination number for each grain in the specimen COEA01, through steps 01 to 17. In this figure, two distinct families of histograms can be seen: the histograms for increments at the beginning at the end of the test the histograms for the end of the test. There appear to be some configurations that have histograms at the interface between these two families. The values visible in the histogram plotted in blue represent COEA01-01, *i.e.*, the initial configuration under confining pressure. The mean of the distribution in the initial state of COEA01 has a mean of 8.90 contacts per grain, with some grains having more than 20 contacts (the maximum recorded coordination number belongs to a grain with 25 contacts). These numbers are geometrically impossible with monodisperse spheres, however it is important to recall that even though the Caicos ooid grains are rounded, they are not spherical, and also not monodisperse. The few grains with surprisingly high coordination numbers have been individually checked, and correspond to particularly large grains, often surrounded by small grains.

Figure 6.12 shows the development of the mean coordination number measured both over the whole specimen (black line), and also exclusively in the zone of the shear band for test COEA01 (blue line) against axial shortening. The local measurement of coordination number appears to start from a higher value than the global measurement, which may be explained by the fact that this zone in the specimen has some grains with high contact numbers, and also includes fewer grains on the boundaries of the specimen. The fact that there are more small grains on the outside of the specimen is likely to be caused by the deposition process (*i.e.*, dry pluviation). Normalising the coordination number by grain volume causes these two graphs to start from the same value in configuration COEA01-01. As the test develops both the global and local coordination numbers decrease, both tending towards final plateaus. The global coordination number tends to a value of 7.3 contacts per grain. The local measurement inside the shear band is more noisy in the last configurations, but the plateau appears to be at around 6.0 contacts. Both lines have the

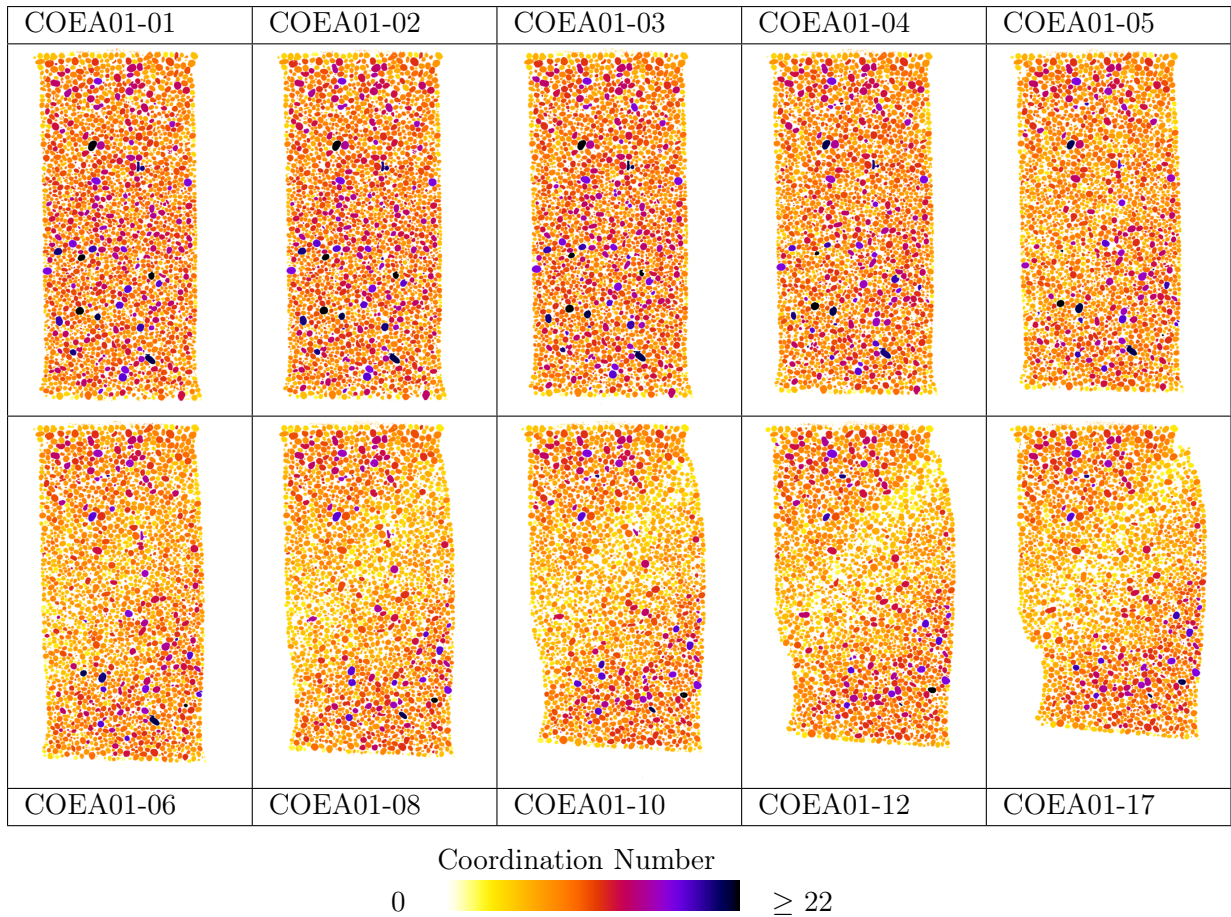


Figure 6.10: Slices from some selected configurations in the test COEA01. Each grain is coloured by its coordination number in that state.

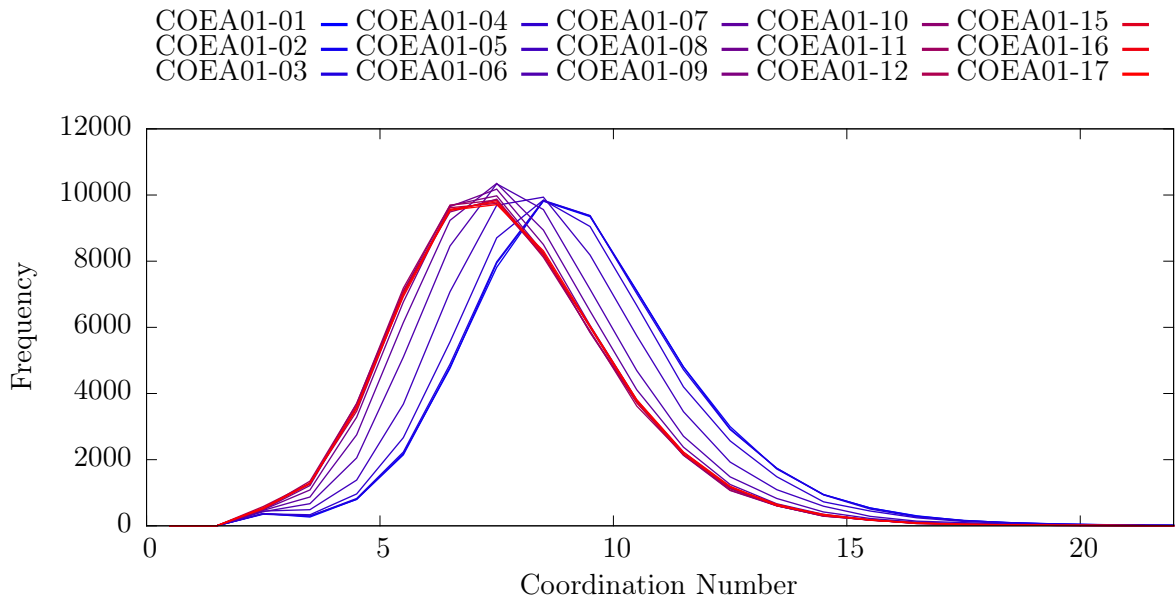


Figure 6.11: Figure showing the different coordination number histograms for selected configurations of test COEA01. The first imaged configuration COEA01-01 is blue and subsequent configurations get progressively red

greatest reduction in coordination number at around 3.5% axial shortening, which corresponds very closely to the peak in axial stress of the specimen.

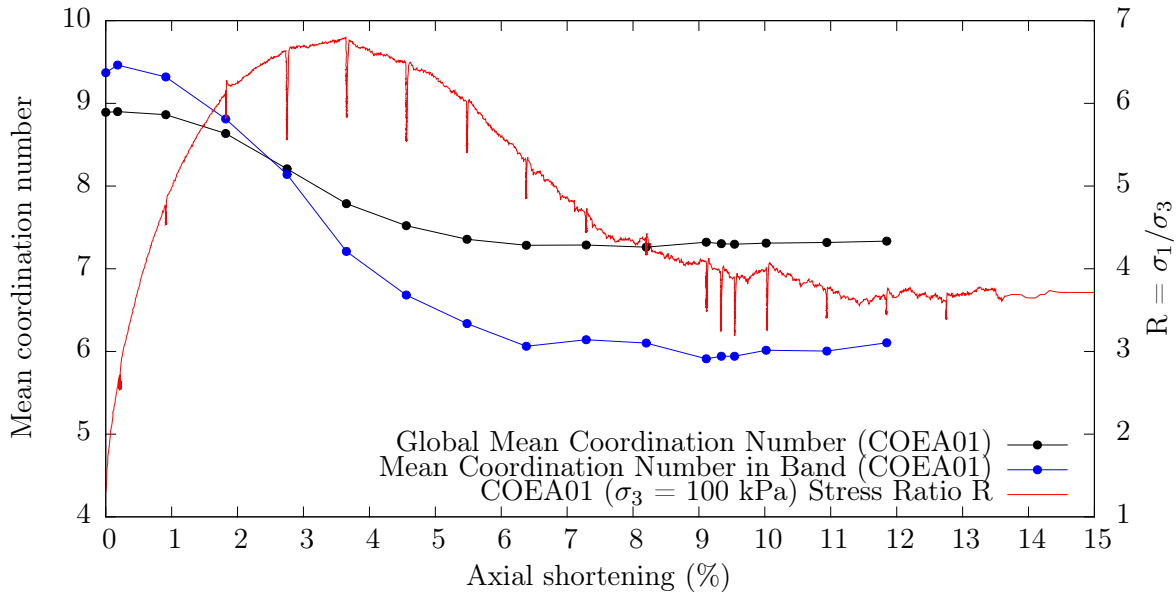


Figure 6.12: Figure showing evolution of the mean coordination number of the entire specimen (black line) and in the shear band (blue line) in all the imaged configurations of test COEA01 against axial shortening. The specimen’s stress ratio is given in red, also against axial shortening

It is interesting to note that the globally-averaged measurement on individual grains seems to be very sensitive to the micro-mechanical events occurring in the specimen. It appears that measurements coming from contact statistics are very sensitive indicators of micro-structural change as opposed to measurements based on specimen volume (be it at a global or local scale). If the trends in Figure 6.12 are compared to the evolution of the volumetric strain in Figure 6.2, or porosity in Figure 6.6, it is interesting to see that at the end of the test COEA01 when the residual stress state has been reached, the volumetric strain keeps increasing from configurations COEA01-11 onwards, whereas the mean coordination number becomes constant at around configuration COEA01-10. It appears as though the measurement of the average coordination number (even at the scale of the specimen) responds more rapidly to change in the way in which the force is distributed inside the specimen, than the volumetric measurements.

The performance of the contact based measurements is perhaps not surprising for two reasons, first of all contact based measurements are much more sensitive to small grain re-arrangements, whereas volume-based measurements require larger grain re-arrangements to be able to detect changes. Furthermore, contacts are the points at which forces are transmitted through the granular skeleton, which means changes in the way in which the force is carried at a meso-scale will be reflected rapidly in measurements coming from contacts.

The difference between global and local measurement of coordination number measured in these tests confirms the observations made in 2D in Sibille and Froiio (2007), which are reproduced in Figure 6.13. The coordination numbers cannot be compared, for reasons both of grain shape, but also because of the number of dimensions involved; the general trends however are very comparable to the observations made herein.




Figure 6.13: Figure taken from Sibille and Froiio (2007). Original caption: Evolution of the coordination number during the biaxial compression phase, for the whole specimen and inside a shear band

6.3 COEA01: following processes in 3D

This section highlights some of the measurements that it is possible to make between increments, *i.e.*, analysis of kinematics using the tools described in Chapter 5. The main two tools for this are the Continuum DIC code TOMOWARP, and the discrete grain tracking technique, developed in this work, called ID-Track.

TOMOWARP makes measurements at regularly spaced intervals, scalar components of which can be conveniently visualised as 3D images. ID-Track however calculates information on a per grain basis. When this needs to be visualised in 3D space, slices of a labelled volume are extracted (as stated in the introduction to this Chapter, slice will be oriented in such a way to include both the axis of the specimen, as well as the normal to the shear band that develops by the end of each test – unless stated otherwise). These slices are then used to generate plots whereby each grain label is replaced by a scalar value that has been calculated for that given grain. These slices are shown with a colour map that translates scalar values to a colour for easier visualisation in 2D print.

This section starts by showing the kinematic fields that result directly from the application of ID-Track. The most immediate output of ID-Track is the per-particle measurement of displacement. Results of grain rotations shown come from the the hybrid Digital Image Correlation and ID-Track approach developed in this work. The concepts of strain at the micro-scale are investigated by comparing the strain field obtained by two different methods: the first consists in using the discrete grain-based kinematics combined with a definition of neighbours to each grain based on triangulation of grain centres in order to calculate the derivative of displacement at the local scale. The second method is simply the application of the Continuum-DIC code TOMOWARP to the greyscale images coming from tomography.

6.3.1 Grain displacements

ID-Track measures and saves a three-component displacement vector for each grain (in the coordinate system of the image), so any scalar component of this vector can be shown. Furthermore, changes of base (for example into a coordinate system aligned with the shear band) are trivial to compute, as will be illustrated at the end of this section.

Figure 6.14 shows grains in configuration COEA01-01, coloured by their measured displacement over increment COEA01-01-02. Grains at the top of the specimen have zero displacement, since this is imposed by the triaxial system (loading is applied from the bottom). It can be seen that, in this increment, the field of grain displacements is axially symmetric (at least in this slice) and increases smoothly from 0 pixels at the top of the specimen, to approximately 3.3 pixels at the bottom of the specimen. Some untracked grains (shown in grey) are visible in this slice, but they do not inhibit the interpretation of the results. The increment COEA01-01-02 shown is a very small increment at the beginning of the test, where the specimen has been deformed by a very small amount. The displacement applied is so small that it is close to the sensitivity of the LVDT – the point at which loading is stopped is measured at being between 0.0491 and 0.0556 mm by LVDT, with the average being 0.0518 mm (*i.e.*, 51.8 μm). This corresponds well to the displacement measured for the grains at the bottom of the specimen (51.6 μm).

The smoothness of the distribution of the vertical displacements shows that the specimen is deforming homogeneously and that the measurement of displacements is highly accurate: at the bottom of the specimen the imposed loading of around 50 μm corresponds to 3 pixels (a technique able to measure displacements to the nearest pixel would only have identified three bands of different values in this measurement).

Despite representing only a 0.22% shortening the axial stress response the specimen goes from a stress ratio of 1 to 2.65 – a stiff initial response which confirms that the specimen has a satisfactory initial geometry.

COEA01-01-02 vertical displacement
(Negative is upwards)




Figure 6.14: Vertical slice of configuration COEA01-01 extracted to contain the axis of the specimen as well as the normal the the shear band that eventually forms. Grains are coloured by their vertical displacement measured over the increment COEA01-01-02, and untracked grains are shown in light grey (many untracked grains are identifiable as separation errors in this slice)

Figure 6.15 shows the same information as Figure 6.14, but calculated on increments throughout the test. The same lookup table is used for all images in Figure 6.15, which means that the colours representing the values for increment COEA01-01-02 are very low on the scale. For conciseness not all increments are shown – configurations COEA01-14 and -15 are omitted and grains are tracked directly over COEA01-12-15. Increments after COEA01-01-02 show a progressive deviation of the vertical displacement field away from axial symmetry. As early as COEA01-02-03 a slight tilt in the field can be noticed, however the gradient of the vertical displacement field (*i.e.*, the strain in the vertical direction) is almost constant in the direction of the height of the specimen. The gradient of the displacement field however becomes increasingly inhomogeneous as the test develops. By the end of the test, in increment COEA01-16-17, the vertical displacement field is composed of two bodies of sand undergoing rigid body displacements with a zone a few grains wide where there is a large displacement gradient. The horizontal profile visible in the lower body of grains is consistent with a rigid-body rotation (due to the ability of the bottom platen to rotate). This sharp change in the vertical displacement field is visually similar in all increments until COEA01-08-09 *i.e.*, on the downwards slope of the strain softening regime.

Between COEA01-02-03 and COEA01-08-09 two key developments can be seen to take place, the first is that the gradient becomes increasingly large and – most importantly – localised. Second, the inclination of the contours of equal vertical displacement can be seen to tilt increasingly from horizontal until reaching a final value of around 47° from horizontal.

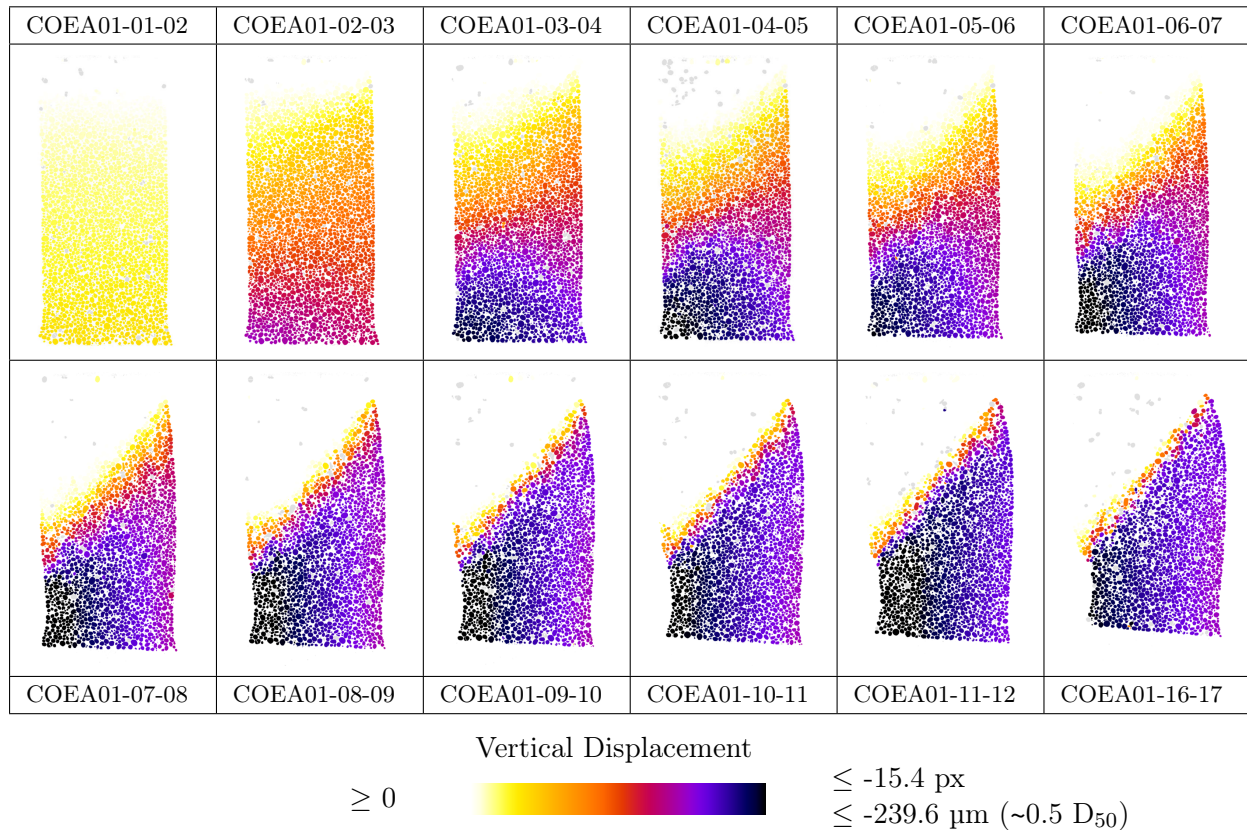


Figure 6.15: Vertical slices of specimen COEA01 showing grain kinematics measured over increments during the test. Grains are coloured by the value of their measured vertical displacement over an increment (negative values mean upwards), in the initial configuration of the increment. Untracked grains are shown in light grey. The same lookup table from scalar to colour is used for all images. Not all increments of the test are shown

In order to stress the 3D nature of the measurements, the same procedure used to create the slices shown in Figure 6.15 (*i.e.*, colouring of grains by a scalar quantity) has been applied to a 3D image for the vertical displacements measured over increment COEA01-16-17; this is shown in Figure 6.16.

The observation of the progressive development of the vertical displacement field is underlined by Figure 6.17. This figure shows the displacement profile along the axis of the specimen in the various increments studied. The first increment (COEA01-01-02) is blue, and subsequent increments tend increasingly towards red. With the exception of a poorly tracked grain in increment COEA01-12-15 all the displacement profiles start from 0 displacement at the top of the specimen. As the test progresses, the last points along the vertical profile get progressively closer to the top: this is expected, since the specimen is being shortened. Increment COEA01-01-02 (blue) has a profile of vertical displacement with a constant gradient, decreasing from 0 at the top of the specimen to approximately -3 at the bottom, as confirmed above. The following increment, COEA01-02-03 has a larger displacement applied at the bottom of the specimen, however the gradient of the vertical displacement field along this profile remains constant. Subsequent increments have been regularly spaced in terms of axial displacement, which explains why they all achieve a final value of around -14 pixels vertical displacement. Increment COEA01-03-04 is approximately well described by a constant gradient, thereafter however, the gradient (*i.e.*, strain




Figure 6.16: 3D rendering of the entire COEA01-16 with grains coloured by their measured vertical displacement in increment COEA01-16-17

the vertical direction) becomes most intense in the middle of the specimen. The profiles in the bottom of the specimen can be seen to become increasingly horizontal until increment COEA01-08-09, which is horizontal and similar to subsequent increments. This confirms the observation that the localisation of the displacement gradient is fully developed by stage COEA01-08-09 and remains constant thereafter. Important observations possible from this figure are that firstly the gradient in the vertical displacement field becomes large in the middle of the specimen, but without becoming discontinuous – showing clearly that the region in which strain is localised becomes gradually thinner as the specimen shears, however it has a finite thickness. Furthermore, this grain-wise (and therefore constant-by-parts) observation of the vertical displacement gradient seems to show a variation of the rate-of-change of the gradient, *i.e.*, the second gradient, without it being discontinuous.

Change of base of displacements

The 3D displacement vector calculated for each grain, over each increment can easily be projected into another base. A particularly relevant and useful change-of-base which can be performed is into coordinates based on the shear band. Figure 6.18 shows an orthonormal base (a, b, c) based around the orientation of the shear band, where component a is the unit vector in the shear band which is in the direction of maximum slope. Component c of this base, is the unit vector which is normal to the shear band. The last component, b is orthogonal to both a and c and is consequently the horizontal vector in the plane of the shear band.

The slices in Figure 6.18 show components of grain displacement in increment COEA01-16-17 in a base aligned with the shear band. Slices showing components of displacement in the a direction reveal that outside the shear band there is very little displacement in the a -direction

COEA01-01-02 – COEA01-04-05 – COEA01-07-08 – COEA01-10-11 – COEA01-15-16 –
 COEA01-02-03 – COEA01-05-06 – COEA01-08-09 – COEA01-11-12 – COEA01-16-17 –
 COEA01-03-04 – COEA01-06-07 – COEA01-09-10 – COEA01-12-15 –




Figure 6.17: Vertical profile of incremental vertical displacements measured along the axis of specimen COEA01. The first increment of the test (COEA01-01-02) is blue, and subsequent profiles become increasingly red

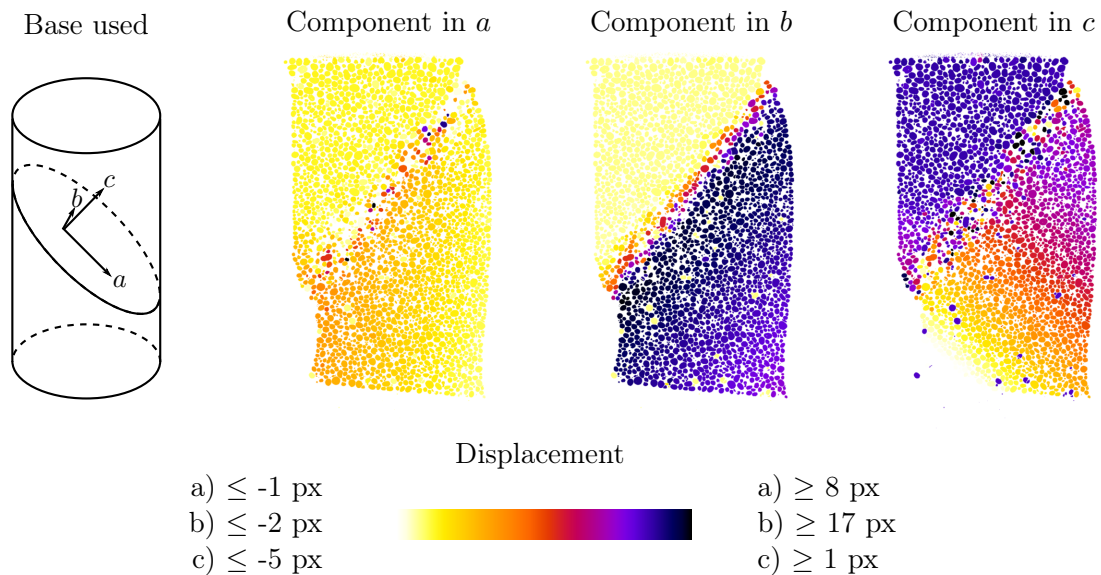


Figure 6.18: Slices showing the components of the 3D displacement calculated for each grain in increment COEA01-16-17, projected into a base based on the orientation of the shear band

(*i.e.*, normal to the direction of the slices), although the left part of the bottom block of grain does show a very small gradient, revealing that there is a small component of rotation in the bottom block of grains. Inside the shear band, displacements in the a direction show some considerable inhomogeneities, with some considerable grain motion in the direction of a .

The slice showing displacements in the b -direction shows the bottom block of grains displacing in the direction of the shear band. In the direction of c the displacements in the b -direction slightly reduce in the lower block; this is consistent with the rotation of this block of grains with the rotation of the lower platen – this can also be confirmed in the slice showing components of displacement in the c -direction, the top part of the lower block of grains has little displacement

in this direction, whereas the lower part is moving in the c -direction. Again, the displacements in the c -direction reveal several grains inside the shear band which have displacements considerably different from their neighbours.

6.3.2 Grain rotations

The kinematics of grains (which are assumed to be rigid bodies, at the relatively low mean stresses that specimens are subjected to in this work) are composed of displacements and rotations. Grain displacements are a relatively natural output of ID-Track, whereas rotations are more challenging to measure. Section 5.3.6 shows a grain-based DIC approach which has been developed in this doctoral work order to measure grain rotations. This tool considerably improves the measurement of rotations from the previously implemented technique (which orients grains based on their moment of inertia tensor). Following the discussion in Section 5.3.7, the grain-based DIC technique for measuring rotations will be used to make measurements in the rest of the Chapter.

Figure 6.19 shows the same slices as shown for displacements in Figure 6.15 but with grains coloured by their total scalar value of rotation. The fields of rotation values at increments at the beginning and at the end of the test (*e.g.*, COEA01-01-02 and COEA01-16-17) show two very different mechanisms. At the beginning of the test, grain displacements are shown to correspond to a relatively homogeneous deformation gradient with what is expected to be uniform strain at the specimen scale; limited amounts of particle re-arrangement are expected to cause small amounts of grain rotation, which increment COEA01-01-02 confirms. By the end of the test the shear band is fully developed, with very localised strain accompanied by intense rotations of a band of grains around 8 grains wide.

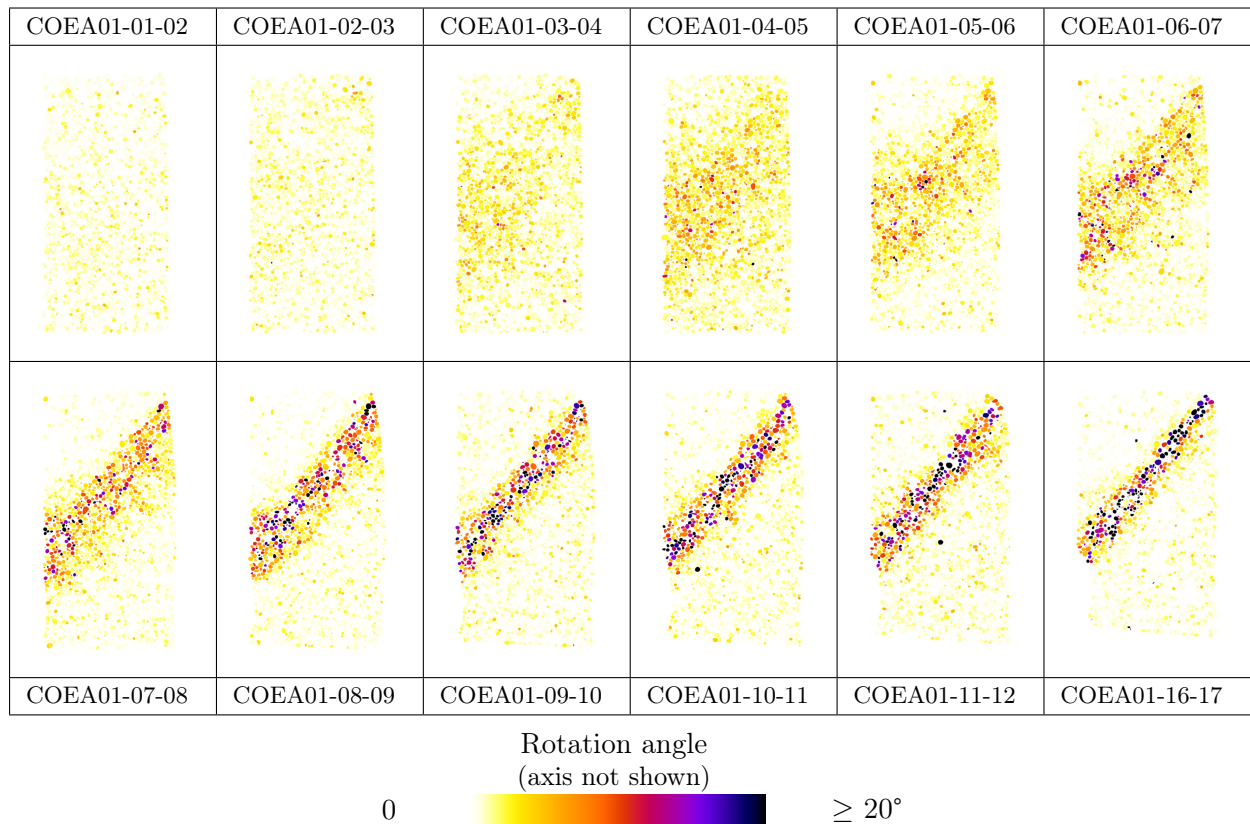


Figure 6.19: Slices of test COEA01 showing grains coloured by the angle of rotation from the axis and angle representation of rotations measured on several increments. Grains are coloured in a slice corresponding to the beginning of the increment. Not all increments are shown

The transition between these two modes of deformation is considerably different than the same transition viewed and measured through grain displacements. The measurement of rotations in increment COEA01-02-03 shows a little localised activity in the top right hand corner of the specimen, but not a significant deviation from the uniform field of close-to-zero rotations measured in the first increment of the test. Over increment COEA01-03-04 a structure around 13 grains wide is clearly visible crossing the specimen. Until increment COEA01-06-07 (*i.e.*, up to the peak in the specimen’s stress response) the width of this structure does not appear to reduce, however this band of rotating grains appears to progressively develop lines of rotating grains which are conjugate to the shear band in the lower part of the specimen. In increment in increment COEA01-06-07, five different “fingers” can be identified in the lower half of the specimen, traces of some of these can be seen in preceding increments, and can be seem to persist into the next increment. These fingers appear to be composed of lines of single grains rotating with relatively similar values of rotation. Fingers of rotating grains appear to die out from increment COEA01-08-09 onwards, with rotations concentrating in the shear band.

Besides the fingers in the lower part of the specimen, in increment COEA01-06-07 networks of rotating grains organised in the direction of the shear band are also visible in this slice. Two parallel lines of rotating grains are particularly visible in the top part of COEA01-06-07, with some indications of these parallel lines in the preceding increment. These two parallel lines of rotating grains precisely delimit the final position of the shear band that develops by the end of the test. In these two slices grains which have large values of rotation are often accompanied by other grains which also rotate. These could be interpreted as rotation centres where a single grains which has experienced a relatively large rotation has caused grains around also to rotate – this can be seen for example towards the middle of the slice labelled COEA01-05-06.

Compared to the gradual concentration of the gradient of the displacement field shown above, the rotations appear to become inhomogeneous at around the same point in the test (COEA01-03-04), however the way in which the concentration is occurring appears to be different. The initial thickness of the band of grains with increased rotations appears to be well defined from the beginning of its development. Furthermore, the angle of the band from horizontal also appears to be well defined, and does not change throughout the test, unlike the iso-levels, or contours of the displacement field. Some structures with strong alignment parallel and conjugate to the shear band are seen to develop and then reduce as the shear band matures and concentrates. The displacements measured appear to concentrate from a wider distribution than the rotations, and to concentrate more rapidly. No geometrical structures are visible in the data as presented in Figure 6.15, but is perhaps because this structure is hidden in the data since a wide range of values has to be plotted for the displacements. Small deviations away from a local value of displacement would allow these structures to be picked out from the displacement data; this is precisely what is done in the following section, which shows micro-strain as calculated on grain displacements (and compares it to strain measured from continuum digital image correlation).

6.3.3 Axes of rotation

The axes of rotation of the grains in the shear band are intuitively expected to be aligned: in the slices shown, grains are expected to rotate around an axis normal to the slice (a in Figure 6.18), *i.e.*, around an axis perpendicular to both the major principal stress and the normal to the shear band. The axes of rotations of the grains have not been shown in Figure 6.19: grains are simply coloured by the measured value of angle of rotation.

In order to confirm the polarisation of the axes of rotation of grains in the shear band, all the grains within the zone of the shear band in configuration COEA01-16 are extracted, and their axes of rotation plotted for rotations measured on the increment COEA01-16-17. 6720 grains have thus been selected, and their axis of rotation plotted as stereoplots in Figure 6.20.

A small amount of alignment can be seen in this plot, at the top and bottom of the circle, representing the orientation expected for grain rotations, however the axes of rotation are well-




Figure 6.20: Orientation of rotation axes of grains inside the shear band over increment COEA01-16-17. The stereoplot on the left shows the raw data, the plot on the right preserves the same projection used on the left, but “bins” the data into small sectors, shows the number of points per bin, normalised by median number of points per bin. The yellow dot indicates the orientation of the *direction of maximum slope in the shear band*

distributed in all orientations.

6.3.4 Calculation of local strain

This section deals with the measurement of strain at the particle scale. First a grain-based method will be shown for the calculation of strain, which will then be compared to results from continuum DIC. Although strain is the output of both methods, there is a fundamental philosophical difference in the way that strains are both obtained and defined which it is important to underline. In the case of continuum DIC, the specimen’s discrete nature is ignored, and the kinematics of the deforming material (from which strain is then derived) are measured at regularly spaced intervals. The grain based method is fundamentally different: the kinematics of the deforming medium come from the displacements measured from the displacement of grain centres. Grain centres are then triangulated, and strain defined piece-wise on the tetrahedra which span the entire domain of the specimen.

Figure 6.21 shows vertical slices through the field of the deviatoric strain, as computed by Bagi’s method (Bagi, 1996) implemented in the YADE (Šmilauer *et al.*, 2010) discrete element code (into which the grain’s incremental displacements measured by ID-Track are used as an input – the simulation part of the code is not used). Strain is calculated on the tetrahedra which are used to triangulate the grain centres. This field that is visualised is therefore a slice through a triangulated volume, with strain defined on tetrahedra. The strain tensor for each tetrahedron is recovered from the output from YADE and its second invariant is subsequently computed as follows:

$$MaxS = \frac{1}{3}\sqrt{2(\epsilon_{XX} - \epsilon_{YY})^2 + 2(\epsilon_{XX} - \epsilon_{ZZ})^2 + 2(\epsilon_{YY} - \epsilon_{ZZ})^2 + 3\epsilon_{XY}^2 + 3\epsilon_{XZ}^2 + 3\epsilon_{YZ}^2} \quad (6.1)$$

This non-negative quantity is referred to as deviatoric strain, and (as shown in Equation 6.1) is four thirds of the radius the of the Mohr's circle of strain (*i.e.*, maximum shear strain) in triaxial conditions.




Figure 6.21: Deviatoric strain calculated from grain kinematics

In section 6.3.1 it was observed that the gradient of vertical displacements of grains appears to concentrate from a wider area of localised behaviour, and stabilises more rapidly than grain rotations. Figure 6.21 confirms this: some preliminary, widely distributed signs of localised strain appear as early as COEA01-02-03 with a wide characteristic thickness, and these concentrate into a thin band which achieves its final thickness around increment COEA01-09-10. This should be compared to the measurement of rotations, which starts from a well defined thickness, which persists longer into the test, meaning that the grain rotations concentrate in a later increment. The final width of the band of concentrated rotations appears to be larger than the final band seen here, although individual grains cannot be seen in this visualisation.

These fields also allow the structure identified in the rotations of grains in increments COEA01-05-06 and COEA01-06-07 (such as the double two parallel bands of activity, as well as the “fingers” starting from the shear band in the bottom part of the specimen) to be confirmed.

In the increment just after the peak (COEA01-07-08) the mechanism of shear strain concentration appears to change. Whereas up to this point a band of more or less constant shear strain has progressively got thinner and more intense, in increment COEA01-07-08 a second phase of concentration can be seen to operate: a zone of considerably increased shear strain is visible *inside* the shear band. This is the point at which the width of the band of concentrated rotations starts to reduce, and corresponds to an increment soon after peak, in strain softening. This zone of intense shear propagates then through the specimen. In the final increments of the test, however, the maximum strain measured along the shear band is not constant. Blocks that experience large amounts of shear can be seen to be interspersed by blocks which shear less.

In summary, the concentration of the shear strain in this specimen appears in two distinct, but overlapping phases: a first phase is a concentration *in space* of shear, in which the width of the band of increased shear decreases while having a constant value of shear strain inside; this effect can be seen to occur principally up to the peak stress. The second phase is concentration of the intensity of shear strain within the band, which appears after peak, and which propagates through the specimen, finally causing a reduction in thickness of the band.

It can be noticed in increments COEA01-06-07 (bottom of the specimen), COEA01-08-09 (bottom) and in increment COEA01-11-12 (top) that just one incorrectly tracked grain has a considerable effect on the shear strain maps produced. This in itself confirms that most grains are tracked successfully.

Figure 6.22 shows slices of deviatoric strain as calculated by TOMOWARP. The measurements outside the specimen (because digital image correlation in this case is performed on a regular grid) have been removed by hand from these images. Inside the specimen, the field of measured deviatoric strain compares well to the fields measured using particle kinematics. Furthermore, the continuum DIC has been performed with $49 \times 49 \times 79$ calculation nodes, meaning there is in the order of one node per grain in the specimen (although calculation nodes are not aligned with grains). The slices extracted from this volume are consequently images with 49×79 pixels – which is why they appear pixelated in Figure 6.22




Figure 6.22: Deviatoric strain measured by TOMOWARP

The output from TOMOWARP appears to be more noisy inside the specimen, this is likely because the sub-pixel part of the image correlation algorithm has some preferential values (at 0.5 pixels displacement in the case of TOMOWARP). The bias towards preferential values has little influence in the measurement of displacements, however in gradients are computed from these fields, the noise they cause can be seen. Measurements based on particle kinematics are much less vulnerable to this because positions are calculated using a particle's centre of mass (which is calculated with an error of less than 0.05 pixels, without an apparent bias in the measurement).

Particle based tracking is fast, once a 3D image is separated and labelled. In this work,

ID-Track has been combined with a relatively optimised commercial image processing package (Visilog[®]) which is able to separate and label a volume in around two hours. The calculation of local strain with YADE takes only a few minutes – this compares very favourably to the processing times of 122 hours for the $49 \times 49 \times 79$ volume analysed with TOMOWARP (both on the same sort of machine). The images coming from the triangulation of grain displacements are defined vectorially, and therefore do not suffer from the same sort of pixelisation.

6.4 COEA01: discussion of processes observed

This section has shown how the tools developed and described in the previous Chapter can be used to analyse a single test. These tools allow many insights to be gained into the grain-scale processes behind the localisation of strain in this specimen.

The measurement of single images at the meso-scale shows that the overall dilation of the specimen seen in the volumetric strain is due to a large increase of porosity concentrated in a band, with little change in porosity elsewhere in the specimen (see Figure 6.3). The band is measured as being around 11 grains wide by the end of the test. The evolution of specimen's global porosity is compared to the evolution of the porosity in the band (in Figure 6.6); this reveals that both local and global values start at around the same overall porosity (35%), and both stabilise by the end of the test – the global value at 39% porosity, whereas the porosity in the shear bands stabilises at around 47%. Scrutiny of the specimen's initial state of porosity before deviatoric loading does not reveal any significant higher porosity which could explain the position of the band of localised porosity that develops by the end of the test.

Analysis at the grain-scale reveals that the grain-size distribution can be measured reliably for this material, the distribution of grain sizes in the first and last images is very close, indicating that no grain breakage is taking place. The grain size distribution coming from the 3D images (based on the diameter of the equivalent sphere of each grain in the specimen) compares moderately well to the grain size distribution from sieving (supplied by ExxonMobil) for this material, the D_{50} values are however quite far apart (420 μm from sieving and around 325 μm from tomography) as shown in Figure 6.8. A sieve analysis performed in this doctoral work only help to prove that sieve analyses are not easy to reproduce systematically. The correspondence between sieve analysis and the mean sphere diameter is better for the other two materials studied in this work.

At an even smaller scale, contacts between grains are characterised in the different 3D images of specimen COEA01, showing that overall the contact numbers reduce, and that contact numbers particularly reduce in the zone where porosity increases (shown in Figure 6.10); the correspondence between these two quantities is expected.

The specimen starts with a mean coordination number of 8.9 which reduces to 7.3, whereas the coordination number in the shear band reduces to 6. When comparing the evolution of this measurement to the evolution of porosity, it is visible that the changes in the coordination number are correlated much more closely than porosity to the specimen's macroscopic stress response; for example, the coordination number reaches a plateau as soon as the residual stress is reached, whereas the porosity, or volumetric strain stabilise later in the test. This is possibly explained by the fact that at the grain scale the macroscopic stress corresponds to forces between grains, which are carried by contacts, which in turn can be identified with relative confidence. Changes in porosity are slower to respond because in order to have changes in local porosity grains must displace considerably, especially when compared to the displacement needed for a contact to be lost in the tomographic images.

The grain kinematics measured in this test reveal that the specimen starts from very homogeneous conditions: the vertical displacements of the grains in the first small increment at the beginning of the test are axially symmetric (see Figure 6.14), as would be expected for material points inside a homogeneous cylinder subjected to axial compression. Thereafter the vertical

displacement field tilts and starts to concentrate into a highly discontinuous field. A vertical profile of the field in Figure 6.17 reveals that in the first three increments the vertical profile of the vertical displacements remains relatively linear, thereafter a zone of high gradient of displacement starts to appear in the middle of the specimen. The gradient gets increasingly large, tending towards a given gradient and thickness (*i.e.*, it does not become a discontinuity) as strain localises.

Grain displacements are projected into a base aligned with the shear band (Figure 6.18), revealing that in the shear band there are grains which have some considerable components of displacement in directions away from the direction of progression of the shear band. The fact that grains inside the shear band can be seen to have unequal components of displacement in directions away from the direction of the shear band (*b*) is not surprising: as the shear band develops, strain localises into a thin band of grains which have differential displacements relative to each other. As shown in this section, there appears to be a critical porosity and coordination number that is achieved when the band is fully developed (around 47% and 6 respectively for this test). Although porosity is increased and coordination number reduced, grains are not free to move only in the direction of the development of the shear band, since this would involve a much more significant loss of contacts. Grains therefore rearrange in the shear band, creating enough freedom to shear. Put differently, there is a quasi-2D phenomenon (a shear band) trying to occur in a 3D specimen made of a 3D material. Although the overall displacement vectors of the grains in the lower part of the shear band will be principally in the direction of the shear band, some small variations are expected, as grains move around each other.

The fact that grains do not roll directly over each other is confirmed by the stereoplot of the rotation axes of measured grain rotations in a mature shear band, shown in Figure 6.20. The right, binned part of the figure shows light alignment in the direction of the band, however the plot on the left confirms that the rotation axes are well-distributed in space.

The measurement of grain rotations throughout the test (shown in slices in Figure 6.19) shows some geometrically complex structures in increments just around the peak axial stress. The “fingers” of grains conjugate to the shear band in the lower part of the specimen correspond to the zone of slightly increased porosity of the specimen – the rotation of grains causes local dilation and an increase in porosity. The lack of symmetry in the “fingering” observed may be due to the fact that the specimen’s lower platen is rotating; this fact would also explain why the effect of increased porosity is not symmetric. This opens myriad possibilities for analysis, some of which have been explored (to different depths) in this doctoral work, as well as in two collaborations from chapter 7. al around the band. The geometrically complex structures that appear a little before peak and which die away as the shear band in the specimen concentrates into its final configuration are reminiscent of the “temporary bands” seen in biaxial compression in Finno *et al.* (1997).

These structures perpendicular to the shear band die away as the shear band develops and concentrates towards the end of the test.

Measurements of micro-strain calculated from grain displacements are shown in Figure 6.21, and compared to TOMOWARP in Figure 6.22. The two methods compare well, despite the fact that they are derived using very different techniques. The triangulated strain coming from grain kinematics appears to give smoother results considerably more quickly.

The comparison of the development of the shear band as measured by the deviatoric strain to the development of the band as measured by grain rotations is interesting: rotations appear to concentrate in a well-defined zone before peak and thereafter concentrate slowly, whereas shear strain continuously condenses from a very wide feature in the second increment of the test to a very tight band by the end of the strain softening regime; the rotations take longer to concentrate. The structures seen in the fields of rotations in increments around around the peak, such as lines of grains grains parallel to the band with large rotations, and lines or “fingers” of grains rotating in the lower part of the shear band. The grain-scale comparison between strains

and rotations is hampered by the fact that rotations are defined discretely on grains, whereas strain is defined piece-wise on tetrahedra spanning the specimen. In the residual stress state, some blocks of grains with particularly high shear strain can be identified, whereas in the grain rotations, this same blocky structure cannot be seen. The deviatoric strain and rotation fields overall compare well to each other.

6.5 COEA01: conclusions

The results presented on one test in this section show the micro-mechanical measurements that are possible with the tools developed in this work. The application of these tools to images describing an increment far in the residual stress state allows a characterisation of a mature shear band: its stable porosity and mean coordination values, as well as the width of grains involved in the shear band. Intense grain rotations characterise the shear band, with a “boundary layer” of grains around also having small rotations. The axes of rotations show a little alignment towards grains rolling in the direction of the band, however overall the rotation axes are well-distributed. The shear strain maps of the shear band show some blocky structure, which is not visible in the rotations.

At the beginning of the test, grain kinematics allow the quality of the initial specimen response to be evaluated – in the case of this specimen, the specimen responds very uniformly. This is confirmed by the fact that almost no significant grain rotations or shear strain is measured in the very first increment of the test. However, in subsequent increments, well before peak, the specimen starts to show some non-homogeneous behaviour which eventually develops into a shear band. Analysis of the initial state of the specimen has not been able to provide any clues as to why the specimen chooses to localise where it does.

Around the peak in the specimen’s axial stress response, the shear band is already well-defined although both grain rotations as well as measurements of micro-strain show considerable structure in the non-homogeneous fields measured.

In short, at the peak of this specimen’s axial stress response, some characteristic features of the grain-scale kinematics appear: the localised micro-scale behaviour of the specimen is at the point where it is most organised in space; the increment up to the peak is the last where the initial band of rotating grains has not yet concentrated (*i.e.*, reduced in thickness), and is the first where there is a continuous line of grains with large rotations crossing the specimen.

6.6 Comparative analysis of tests on three different materials

This section will present and discuss one representative specimen from each material and confining pressure from the experimental program carried out in this doctoral work. The emphasis will be in the identification and characterisation of different micro-mechanisms responsible for the differences between the responses of the tests. It is important to recall that in addition to grain angularity varying between the different materials, Caicos ooids are carbonate grains, whereas the other two materials have quartz grains.

6.6.1 Discussion of macroscopic responses

Figure 6.23 shows the macroscopic measurements made on the four tests performed at 100 kPa confining pressure that will be analysed: COEA01, COEA04, OUEA06 and HNEA01. Two Caicos ooids tests are included since there is a considerable discrepancy in the values of the peak observed between these two specimens.

All four specimens behave as expected for dense specimens: they have a clear peak in their axial stress response and are dilatant. The two blue plots, representing the two tests to be

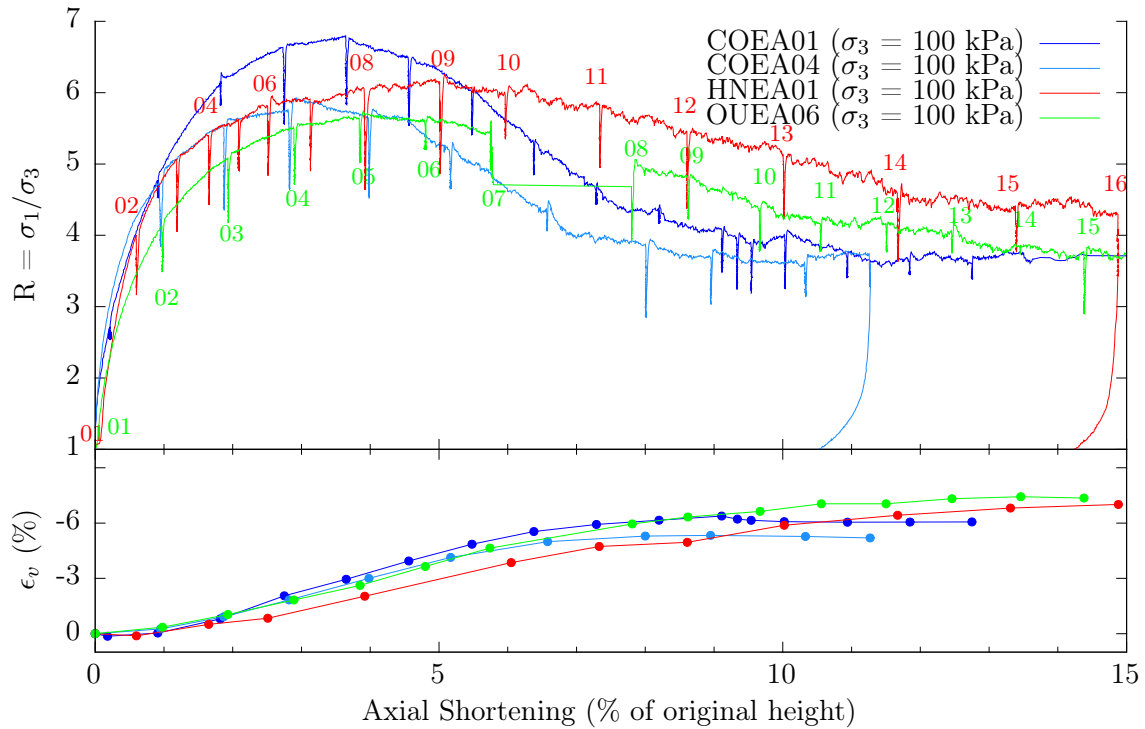


Figure 6.23: Macroscopic specimen responses for COEA01, COEA04, HNEA01 and OUEA06 all tested at a confining pressure of 100 kPa. On the top plot, the stress ratio (R) is plotted against axial shortening, and on the bottom volumetric strain is plotted against axial shortening

analysed on Caicos ooids (COEA01 in dark blue and COEA04 in light blue), have similarly-shaped stress-ratio curves, however COEA01 has a distinctly higher peak stress than COEA04 (a friction angle of 48.0° *vs.* 44.8°), which occurs a little later. Quartz grains with the same degree of roundness as Caicos ooids typically be expected to have a lower friction angle than the other two less rounded materials. However, Caicos ooids are a carbonate sand, and are consequently expected to have a higher peak stress when compared to quartz grains likely due to an increased inter-particle friction. Both tests reach a very similar residual stress ratio, corresponding to a friction angle of 35° . The dilatancy angles at peak for these two tests are very similar; the dilatancy angle of both COEA01 and COEA04 is measured as 20° . Recalling the data presented in Table 3.5, all the external characteristics of specimen COEA04, compared to specimen COEA01, would indicate a specimen that is expected to have a higher axial stress at peak: COEA04 is denser than COEA01, however it is COEA01 that has a considerably higher peak.

The specimen HNEA01 (red lines) of Hostun sand also presents a high peak (with a friction angle of 46°), which appears later than in the other specimens – at around 5% axial shortening. The behaviour of this specimen is dilative like the others, with a dilatancy angle of 17° at peak. It is not clear whether the specimen has reached its residual stress state, however it appears to be close. The stress ratio measured for this specimen at the end of the test is considerably higher than that of the other specimens, stabilising at a stress ratio of 4.4 (friction angle 39°).

The specimen of Ottawa sand (OUEA06) shares the low residual stress ratio of the Caicos specimens and has a peak stress ratio of 44.4° , close to the value reported for specimen COEA04. This peak occurs at 4.4% axial shortening, and is associated with a dilatancy angle of 20° .

The general trends that can be seen in the responses of these specimens are that the rounder materials (Caicos and Ottawa) have very similar values of residual stress, which are low compared to the value of residual stress which appears to have been reached by the angular material (Hostun sand). There is a definite correlation between the amount of axial shortening required to reach

peak and the angularity of the grains in the specimen: the specimens of Caicos ooids reach peak first, followed by Ottawa sand, and last of all Hostun sand. The values of the peak are higher for Hostun than Ottawa sand, as expected. However one specimen of the carbonate Caicos ooids has a very high peak.

Figure 6.24 presents the macroscopic responses of the three tests performed at 300 kPa confinement that will be analysed in this section: COEA02, OUEA02 and HNEA04.




Figure 6.24: Macroscopic specimen responses for COEA02, HNEA04 and OUEA02 all tested at a confining pressure of 300 kPa. On the top plot, the stress ratio (R) is plotted against axial shortening, and on the bottom volumetric strain is plotted against axial shortening

The comparison between the macroscopic specimen responses is very much in correspondence with the comparison with the tests at 100 kPa; Tests COEA02 and OUEA02 have a comparable value of residual stress state, as was found between COEA01, COEA04 and OUEA06 at 100 kPa. The specimen of Hostun sand appears to tend to a higher value of residual stress, consistent with the tests at 100 kPa. As with the tests at 100 kPa, the peak in the axial stress responses is earliest for the most rounded material, and latest for the most angular. Again, the specimen of Caicos ooids has the highest peak, followed by Hostun sand and finally Ottawa sand has the lowest peak – the same trend as in the 100 kPa tests if COEA01 is considered a reliable test. The volumetric response of the specimens also follows the same trends as the test at 100 kPa: the specimen of Hostun sand starts to dilate later than the other two (and this specimen has a longer contractive phase at 300 kPa confinement). Ottawa and Caicos specimens have similar dilatancy angles, with the Ottawa specimen dilating longer into the test. Unlike the tests at 100 kPa, the strain softening part of the macroscopic response of OUEA02 is much more similar to that of COEA02, whereas at 100 kPa the strain softening part of the Ottawa sand specimen follows the response of the Hostun sand more closely.

It is not certain however that the change in Caicos ooids' post-peak behaviour is a material response: in this phase a specimen's response may be strongly affected by its geometry, since even for the two similar backup tests on Caicos ooids performed at 100 kPa the post-peak behaviour is not the same.

Comparing each material's response between 100 kPa and 300 kPa confining pressure shows that the peak stresses occur later in each specimen, and that the values of stress ratio at peak are reduced. Table 6.1 shows the position and value of the peak for all specimens, along with the value of the residual stress.

Specimen	Confining Pressure	Peak R (ϕ')	Shortening at Peak	Dilatancy at Peak	Apparent Residual R (ϕ')
HNEA01	100 kPa	6.15 (46.1°)	5.0%	17°	4.15 (37.7°)
HNEA04	300 kPa	5.35 (43.2°)	7.2%	15°	4.00 (36.9°)
OUEA06	100 kPa	5.50 (43.8°)	4.3%	20°	3.75 (35.4°)
OUEA02	300 kPa	5.20 (42.6°)	5.2%	16°	3.70 (35.1°)
COEA01	100 kPa	6.79 (48.0°)	3.6%	20°	3.70 (35.1°)
COEA04	100 kPa	5.77 (44.8°)	3.4%	20°	3.70 (35.1°)
COEA02	300 kPa	5.79 (44.9°)	5.2%	16°	3.65 (34.7°)

Table 6.1: Table showing some important characteristics of the macroscopic responses of the specimens analysed herein. R is stress ratio σ_1/σ_3

An important conclusion to be drawn from the comparison of all these tests at 100 and 300 kPa confinement is that many of the trends visible between the responses of the different materials at 100 kPa also appear at 300 kPa. These include the order in which specimens reach their peak, the final residual stress states achieved and the volumetric responses of the specimens. However some trends are not the same between the two confining pressures used: the post-peak behaviour of Ottawa sand is closer to Hostun at 100 kPa and closer to Caicos at 300 kPa.

Based on the observations of the macroscopic responses of the specimens discussed above, the principal question that arises is whether the different grain types (angularity and material) are responsible for this behaviour and, if so, whether this is manifested in the grain-scale mechanisms of the specimens.

To give focus to this Chapter, some key questions, based on observations made above, will be addressed:

- Can the stress response of specimens COEA01 and COEA04 be explained through micro-mechanical observations?
- Is the higher residual stress of the Hostun specimens explained by grain angularity?
- Is the relationship between grain roundness and the earliness of each specimen's peak explained by strain localisation occurring in the specimen?
- Can the lower and later response of specimens tested at 300 kPa confinement be explained by micro-mechanical observations?
- Is the different (carbonate) material of the Caicos grains only responsible for the high peak stress?

6.6.2 Comparison of tests COEA01 and COEA04

This section will look for micro-mechanical explanations for the observed difference in the macroscopic responses between specimens COEA01 and COEA04. As stated above, the initial conditions of COEA04 (denser than COEA01) give the expectation of a specimen with a higher peak, but the opposite happens in this case. Figure 6.25 shows the macroscopic responses of both

specimens with the different states that have been analysed labelled for both specimens on the top graph. The dilatancy at peak can be seen to be very close for both specimens; towards the end of the test COEA01 has undergone more dilation however.




Figure 6.25: Macroscopic specimen responses for COEA01 and COEA04. On the top plot, the Stress ratio (R) is plotted against axial shortening, and on the bottom volumetric strain is plotted against axial shortening. All the configurations that have been imaged and analysed are indicated on the top graph

A difference in peak stress could in principle be ascribed to a measurement error, however the fact that both specimens reach the same final residual stress indicates that the measurements are probably trustworthy. The fact that the residual stress is the same for both specimens is a good place to start for the comparison of these two tests.

Figure 6.26 shows some key micro-scale measurements made over a similarly-sized increment in the residual stress state of both specimens (COEA01-16-17 and COEA04-10-11). Figure 6.26 confirms that the specimens, when in the residual stress state (which occurs at the same stress ratio in both tests) also have comparable micro-mechanical activity: although the shear band happens at a different height in the specimen, the grain-scale mechanisms visible in the mature shear band are the same. The vertical displacement fields both show a band 6 grains thick in the middle of the specimen, which has displacements to the two blocks of rigidly-moving grains. The deviatoric strain fields for the two specimens reveal the same phenomenon of a well-defined band, with some patches of concentrated shear within. Grain rotations in the COEA04-10-11 also appear to be distributed in the same way as they are for COEA01-16-17; there is a concentrated band around 5 grains thick experiencing large rotations, and a band around 7 to 8 grains thick experiencing non-zero rotations in both cases. The membrane can be seen to spread out and limit the rotations, on the bottom right of COEA04-10-11, an effect which is less visible in specimen COEA01.

The conclusion to be drawn from Figure 6.26 is that the micro-mechanisms at play in the increments compared appear to be very similar. This similarity can also be found in the macroscopic measurements made in both specimens over these increments: the stress ratio remains constant over both increments, and both specimens have achieved a zero angle of dilatancy confirming that the residual stress state has been reached.




Figure 6.26: Comparison of the micro-scale measurements made in the residual state of tests COEA01 and COEA04




Figure 6.27: Deviatoric strain maps from specimens COEA01 and COEA04 compared for the increments leading up to peak. Please note that the first increment is COEA01-01-03 (*i.e.*, COEA01-02 is skipped) in order to compare axial shortening increment more closely

Looking now at the beginning of the test, the initial increment COEA04-01-02, is practically the same (in terms of axial shortening) as COEA01-01-03, the increments that follow are also relatively well-aligned, the increments to up the peak stress are shown in Figure 6.27.

Figure 6.27 shows slices through the field of deviatoric strain in increments up to (and over in the case of COEA04) the specimen's peak. The first three increments of each test show little differences in the development of the localised shear band; from the second increment in both tests a very widely distributed region of localised shear exists in the specimen, in the orientation of the final shear band that develops. Some small indications of two parallel bands is visible in COEA04-02-03, however the shear strain remains very diffuse. The correspondence of the shear strain maps in the first increment is encouraging, since both specimens' macroscopic stress responses are very similar up to this stage. However in the next two increments the specimens' macro-scale responses start to differ, but this does not appear to be reflected in the deviatoric strain maps. By COEA04-03, the specimen can be seen to "barrel" laterally, which is not seen in COEA01.

In the last increment shown for each test, a considerable change can be seen: a zone of highly concentrated shear strain can be seen starting to develop within the shear band in specimen COEA04, whereas this cannot be seen in COEA01 at this stage, and only manifests itself in increment COEA01-07-08.

The difference in the peaks of specimens COEA01 and COEA04 appears to lie to some extent in the localised micro-mechanics that occur as the specimens are deformed. The fact that COEA04 has its peak earlier than COEA01 can be seen in the fact that the shear band goes into

its second phase of concentration (with intense shear developing inside a concentrated band) in increment COEA04-04-05 (shown in Figure 6.27), whereas this appears later in specimen COEA01. The fact that this second phase occurs earlier in specimen COEA04, explains its reaching peak before COEA01, but may also in part explain the lower peak stress, since the mechanism that leads to failure in this case is already well developed. The fact that no strong difference in microscopic shear strain responses can be seen in the second and third increments of Figure 6.27 means that perhaps the cause of the earlier localisation cannot be found in the kinematics of the specimens.

In conclusion, the micro-mechanisms occurring in both specimens in the residual stress state are very similar, which confirms that the measured responses are likely to be correct and physical. The lower peak in specimen COEA04 appears to be correlated with a considerable difference in the micro-mechanics occurring in the specimen over the peak: the two-step process of shear strain concentration described in COEA01 also occurs in specimen COEA04 and does so earlier in the test, which is likely to be the cause of the lower peak. However leading up to the peak, where the two specimens' stress responses diverge no considerable difference can be detected.

6.6.3 Residual stress states

This section will look at the residual stress states in the specimens analysed in this section. The residual stress state is key for the study of localisation because it is totally governed (particularly in the small, dense specimens which have been tested) by the final shear band. Differences in the value of the residual stress states are most noticeable first between the Hostun specimens and the others, but also for the same specimen at different levels of confining pressure. Since responses differ, study of the residual stress states allows characterisation of this fully localised mechanism of deformation in order to confirm that the tools developed for micro-mechanical analysis are able to pick out the differences between the specimens.

This section will omit COEA04 from the analysis, as its residual stress state has been compared and found similar to COEA01 in the section above. The residual stress states of specimens tested at 100 kPa (COEA01, OUEA06 and HNEA01) will be compared first. The tests at 300 kPa will then also be compared, to see whether changes in the micro-mechanisms of deformation occurring at higher mean stress might explain the fact that residual stresses are systematically slightly lower at higher cell pressure.

Residual stress states at 100 kPa confinement

An increment in the residual stress state will be analysed for each test at 100 kPa: COEA01-16-17, OUEA06-14-15 and HNEA01-15-16. The lengths of these increments are detailed in Table 6.2. COEA01-16-17, OUEA06-14-15 are similarly-sized increments, however HNEA01-15-16 is a much longer increment (almost twice as long as the others), so this must be born in mind when analysing the measurements made on these increments.

Residual Stress Increment	Axial shortening over increment
HNEA01-15-16	1.72%
OUEA06-14-15	0.92%
COEA01-16-17	0.89%

Table 6.2: Table showing the lengths of the residual stress increments for specimens studied at 100 kPa confining pressure

Figure 6.28 shows vertical slices, taken to include the specimen's axis and the normal to the shear band, which is clearly visible in these slices. The shear strain and grain rotation fields

shown for all three specimens confirm that the shear band is fully developed. In the increments studied there is practically no strain or grain rotation outside of the shear band.

In the residual stress states shown, the deviatoric strain can be seen to be strongly localised into a thin band for all specimens. In HNEA01-15-16 the discrete blocks of very intense shear in the direction of the band (identified earlier in COEA01-16-17) can be clearly seen in red. These blocks are situated inside a wider shear band than that in COEA01-16-17 and one which has a less abrupt change in shear strain in the direction normal to the band. This is best viewed by looking at the extent of the values of zero deviatoric strain in these specimens. Furthermore, the most intensely shearing blocks in HNEA01-15-16 are towards the bottom of the band on the left of the slice, and are towards the top on the right side of the slice. No such deviation is present in COEA01-16-17, where the intensely shearing block are in the middle of a tight band.

The deviatoric strain in OUEA06-14-15 is not localised into a band of constant thickness, in the slice shown the left side of the band is thicker than the right side which is tightly concentrated. The change in the shear strains in the direction normal to the band are not as sharp as COEA01-16-17, but more concentrated than HNEA01-15-16.

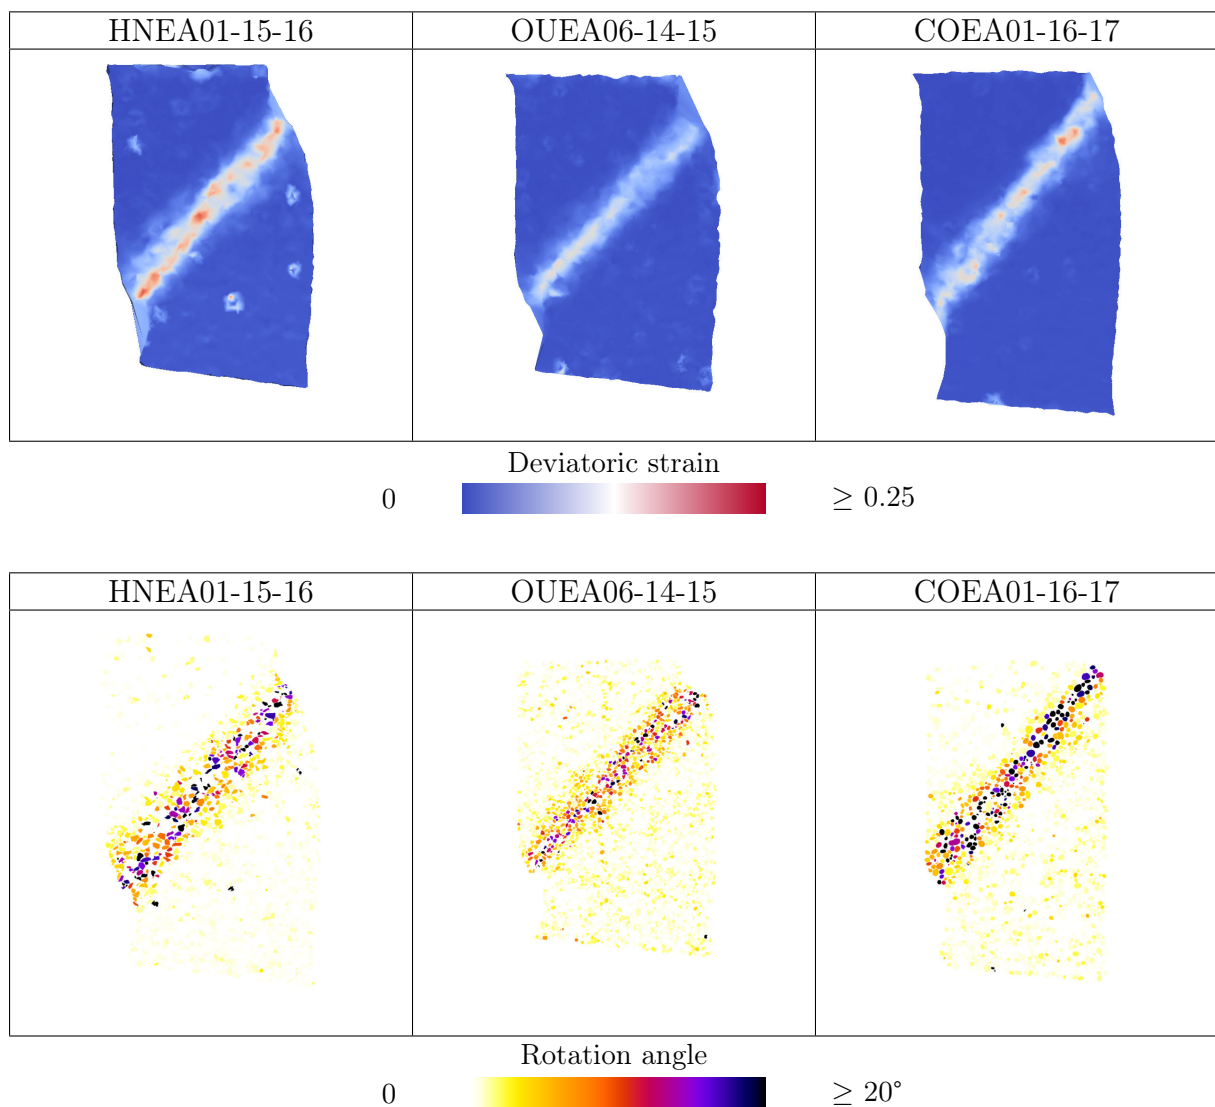


Figure 6.28: Comparison of the micro-scale measurements made in the residual state of tests HNEA01, OUEA06 and COEA01

The mechanism behind the increasing sharpness of shear strains with increasing grain roundness may find its origin in the rotation fields presented in Figure 6.28. In comparing these images

it is important to bear in mind that the grains of Ottawa sand are smaller than the other two specimens, and so have a smaller visual impact than Hostun or Caicos grains in grain-based representations such as the rotation fields. The fields of rotations show a very concentrated band of grains with intense rotations coincident with the band of high shear strain in all three cases. Counting the numbers of grains with intense rotations across the band reveals that six to eight grains have considerably higher rotations than the rest of the specimen in each case. The values of rotations for this band of grains is very large for COEA01, with many grains rotating more than 20° . The values of rotations for both OUEA06 and HNEA01 in their residual stress increments appear to be the same with grains undergoing rotations of around $10\text{-}15^\circ$, but it is important to stress that the increment HNEA01-15-16 is much larger than the other two: if it were the same size as the other two increments studied, much smaller rotations would be expected in the order of $5\text{-}8^\circ$. The intensity of rotations in these three specimens therefore appears to be inversely correlated with grain angularity.

Around the band of intensely rotating grains in increment HNEA01-15-16 grains with smaller, but non-zero, rotations are visible, all along the band on both sides of it. These “secondary rotations” appear to affect a thickness of around three grains on either side of the shear band. The same effect is visible to a smaller extent in OUEA06-14-15, with one or two grains either side of the band of intensely rotating grains having smaller rotations. In COEA01, in many places along the edges of the band there appear to be no “secondary rotations”; several grains with over 20° rotations can be seen to be delimited by grains with close to zero rotation. The secondary rotations are therefore correlated with grain angularity: in the specimens analysed it appears that the most angular Hostun grains have large secondary rotations, which is most likely due to interlocking between the particles. Interlocking is a geometrical phenomenon that causes rotational frustration. Rotational frustration is used in the sense of “inability or difficulty to rotate”, and is expected to affect angular grains considerably more than rounded grains: angular grain shapes are far from spheres, which means that long parts of a grain are able to lie in small pores between other grains. This in turn means that rotational freedom of these grains is considerably reduced and interdependent. In order for one angular grain to rotate when it is packed in a dense granular assembly, grains that prevent its rotation must also be caused to move. Since a shear band has been shown to cause large grain rotations in a thin band of grains, in an angular material the band is wider (less sharp) because the intense core of rotating grains causes grains on the outside of the band to rotate. Figure 6.28 shows that around three layers of grains are involved in these secondary rotations. Furthermore rotational frustration is also likely to occur within the band of rotating grains; if the measurements of rotations within the band were plotted on similar increments of axial shortening, a progressive reduction of grain rotations is strongly expected with increasing grain angularity.

For grains that are far from spheres, the expectation is that rotation induced in the grains on the outside of the band is not fully decoupled from displacement of these grains; which explains why a less-sharply defined band of localised shear strain is visible in the deviatoric strain fields however without the striking nature of the rotation field.

A fully developed shear band in all the specimens tested appears to be composed of a thin band of rotating grains. In Hostun sand, a layer around three grains wide on either side of the band is also caused to rotate by the development of the band; the thickness of this layer is between one and two grains for Ottawa sand, and is further reduced for Caicos ooids. The intensity of rotations inside the band is closely related to the angularity of the grains, with Caicos grains having very intense rotations.

In summary, in the residual stress states of these three specimens different micro-mechanisms are at play: in the rounded material, there is intense rotation in thin band of grains, whereas in the angular material there is a band of grains with less intense rotation that is more widely distributed. In Caicos and Ottawa specimens, when the shear band is fully developed, grain shapes offer little resistance to rotation (and therefore to shear), which is why the band is so

concentrated. In the angular material however, the fact that grains are not able to rotate as freely as in the other materials means that larger resistance to shear is obtained. Although shearing is localised in the specimen, for further shear to occur grains bounding the shear band also need to be rotated (due to grain geometries), which limits the rotations and dissipates energy. In both intensity of grain rotations, and the extent of secondary rotations, OUEA06 is an intermediate case between COEA01 and HNEA01. The reason for OUEA06's similar residual stress state to COEA01 cannot perhaps be read directly from the kinematics of the grains in the shear band; the expectation of the residual stress state given the kinematics is that OUEA06 occupies an intermediate position between COEA01 and HNEA01. However recalling that Caicos ooids are carbonate grains (which explains to some extent the high peak of COEA01), it is likely that the residual stress value of COEA01 is also higher than what would be expected for a quartz material with the same grain morphology than COEA01.

To better understand the micro-scale mechanisms occurring inside the shear band in each specimen, the axes of rotation of the grains that make up the shear band can also be analysed. Figure 6.29 shows the rotation axes for rotations in the residual increments studied for HNEA01 and COEA01. These binned stereoplots are normalised by the *median* value of the number of projected points that enters each bin; the result of this is that bins with a value of 1 (*i.e.*, white in the colour map chosen) represent bins that are at the median value. In both cases, polarisation is visible and, compared to the horizontal orientation of the shear band (shown on the plots), grains can be seen to roll along the slope of the band (with a rotation axis in the direction of a from Figure 6.18). However the polarisation is not very strong. Furthermore, the polarisation of the rotation axes of the grains in the shear band in increment COEA01-16-17 is very specifically in the direction of rolling along the slope of the shear band. In HNEA01-15-16 it appears that the polarisation is both stronger and less well-organised in space than what is visible in COEA01, *i.e.*, there are a larger quantity of grains rotating with an axis of rotation *approximately* in the direction of a .

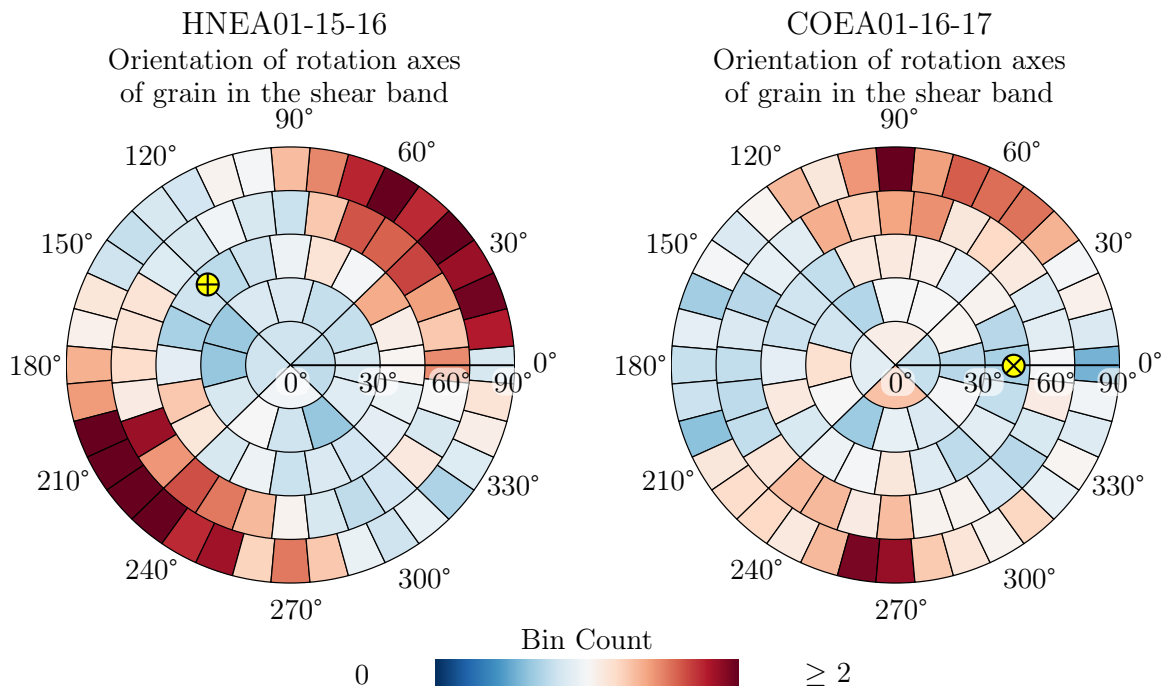


Figure 6.29: Comparison of the rotation axes for the grains in the shear band in specimens HNEA01 and COEA01. The yellow dot indicates the orientation of the *direction of maximum slope in the shear band*

The polarisation of grain rotations gives an unexpected insight into the localisation process

under way: in all cases the rotations are not totally polarised, this has been explained above as being due to the complexity of 3D transformations of rigid particles being involved in a deformation mode that is essentially two-dimensional (a shear band). This said, the rotations in Hostun sand can be seen to be more polarised approximately around a than in COEA01. The fact that the polarisation is a little scattered is likely due to the secondary rotations in the shear band, where rotation is transmitted to neighbouring grains in related orientations. However the increased polarisation of orientations in the angular material is more surprising, however what must be taken into account in this case is that the increment of shortening over which the rotations are measured is larger for HNEA01, which means that in this increment the values are closer to the mean rotation expected.

Figure 6.30 which shows the porosity field for the state at the beginning of all three of these increments. The band of increased porosity in specimen COEA01 can be seen to be more sharply defined (particularly on the top of the band) when compared to HNEA01 and OUEA06, which both have a more gradual progression of porosity from the inside of the band to the outside.

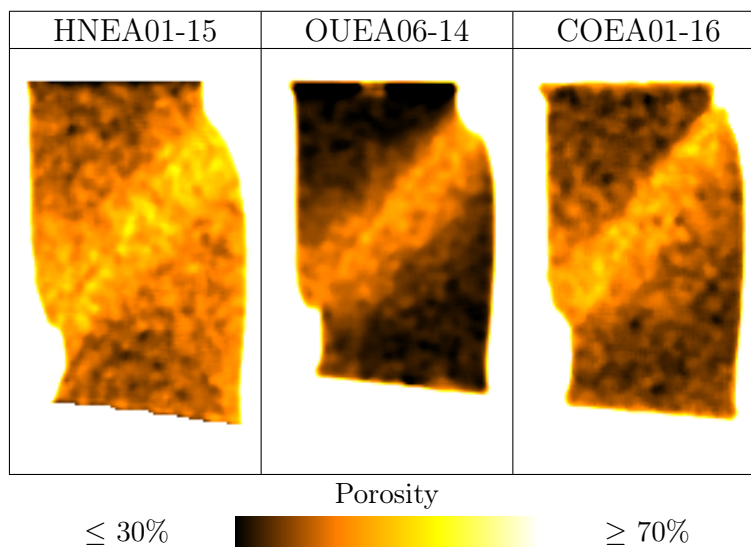


Figure 6.30: Comparison of the fields of porosity measured for states HNEA01-15, OUEA06-14 and COEA01-16

Residual stress states at 300 kPa confinement

This section will look at the residual stress states for the specimens tested at 300 kPa confinement. As recalled in Table 6.1 the residual stress states for the specimens have a tendency to reduce slightly between tests at 100 and 300 kPa. This means that, based on the micro-mechanisms seen above, the bands of localised deformation are expected to be thinner with less developed secondary rotations at higher pressure.

Residual Stress Increment	Axial shortening over increment
HNEA04-14-15	1.86%
OUEA02-11-12	1.36%
COEA02-12-15	1.00%

Table 6.3: Table showing the lengths of the residual stress increments for specimens studied at 300 kPa confining pressure

Table 6.3 shows the sizes of the residual stress increments which will be studied. In this

case all these increment have relatively different sizes, so analysis must be made carefully in these increments. Given the partial polarisation of the rotation axes shown in Figure 6.29, the values of the fields have not been normalised, since it appears that grains do not have monotonic kinematics in the shear band.

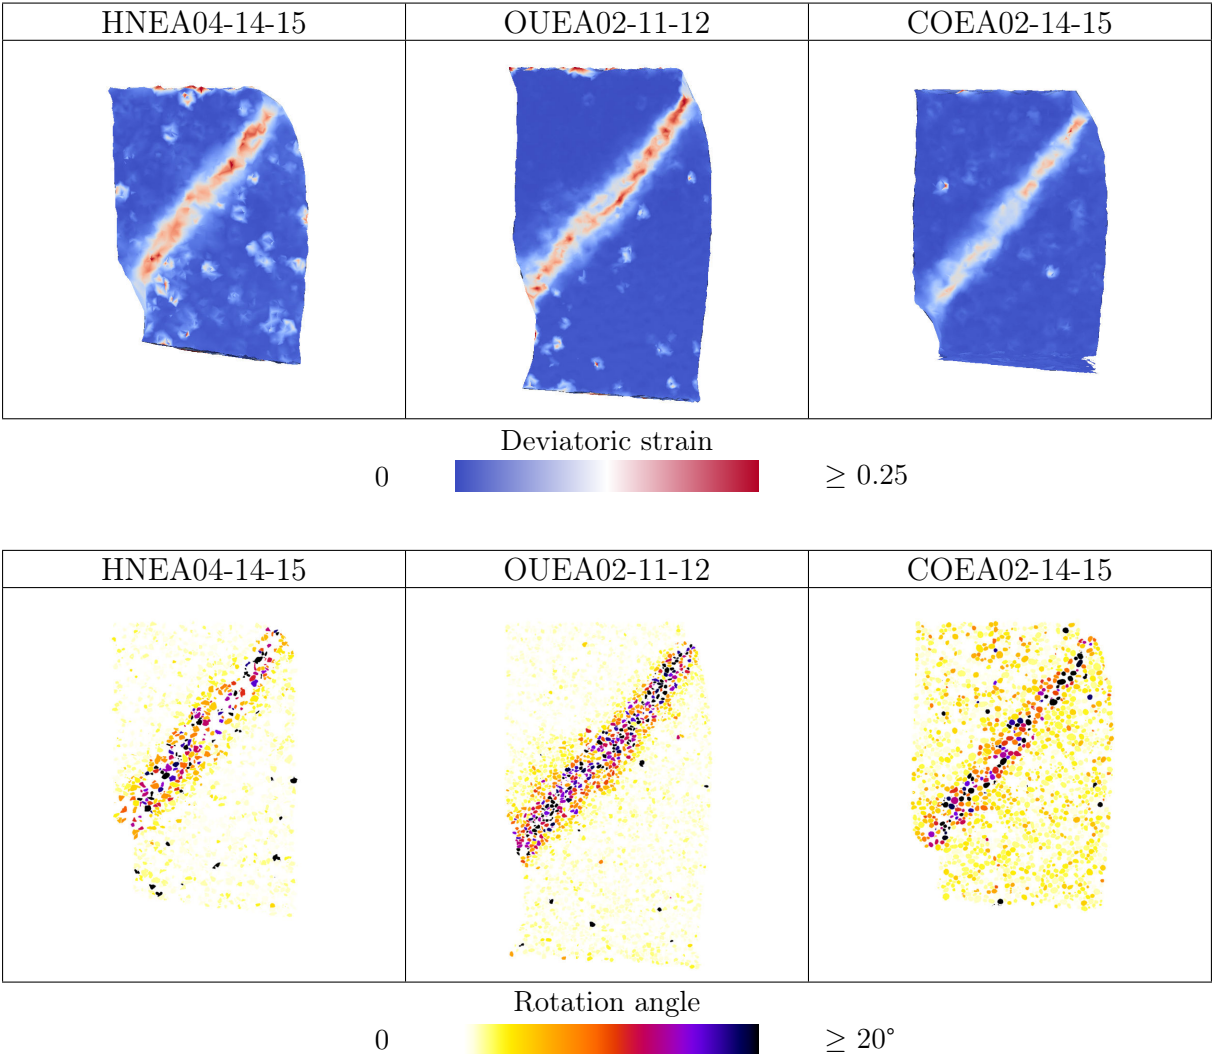


Figure 6.31: Comparison of the micro-scale measurements made in the residual state of tests HNEA04, OUEA02 and COEA02

Figure 6.31 shows the micro-mechanical measurements made on increments in the residual state of three specimens tested at 300 kPa confining pressure. All of the increments analysed are slightly larger than those analysed for the specimens at 100 kPa confinement. A considerable amount of noise is noticeable in the slices from HNEA04 – this is because the combination of the relatively large displacements and the difficulty of reliable segmentation causes some grains to be incorrectly tracked. Furthermore, in the rotation fields for this increment some grains inside the shear band are evidently not tracked.

Comparing the fields in Figure 6.31 to the ones obtained for the same materials at 100 kPa confinement (shown in Figure 6.28) shows a good correspondence. In all the shear strain maps, “blocks” of intense shear can be identified within the shear band. The shear band that has developed in HNEA04 is defined with a smooth border with respect to the material outside, which is not shearing. On the other hand, COEA02 maintains a sharply defined band both in terms of shear strain and grain rotations, which is consistent with the observations at 100 kPa. Specimen OUEA02 also has a band of varying thickness in its residual stress state (as

was observed in OUEA06 at 100 kPa). In this case the band is thinnest towards the top of the specimen (the opposite compared to OUEA06) and, again, the thinnest part of the band is associated with the most intense shear. The rotations of the grains in OUEA02 show that the band thickness increases from around 5 grains at the top of the specimen to around 7 grains at the bottom. The band of rotating grains is much more visible in this increment because the grains undergo larger rotations and consequently stand out more clearly in Figure 6.31.

The extent of secondary rotations in the three specimens shown in Figure 6.31 follows the same trends than the secondary rotations observed in the slices shown at 100 kPa: increasing secondary rotations occur with increasing angularity. Figure 6.32 shows some zooms into the shear band for all three specimens. Although the trend is not striking, reduced secondary rotations can be seen for specimens tested at higher pressure. The challenge with the visual interpretation of these data is that (other than the occasional grain missing in the images of Hostun sand), the values of rotation are different in these different shear bands (due to the different size of the increments). Furthermore, in some cases, the thickness of the shear band in the same material is different: for example the band of rotating grains in COEA02 appears thinner than in COEA01, and conversely, that of OUEA02 appears larger than in OUEA06. However, if the secondary rotations are compared without taking into account the width and intensity of the rotations of grain in the shear band, they can be seen to reduce when confining pressure is increased from 100 kPa to 300 kPa.

Conclusions

A hypothesis explaining the residual strengths of the materials tested is that in the experiments analysed, the two principal factors that affect the residual strength of these materials are the secondary rotations around the shear band, as well as inter-particle friction. The secondary rotations are likely to be important because they are the result of rotational frustration of grains inside the shear band, and for this reason they are linked to the intensity of rotations in the band. The width of the band itself appears however to be more variable (at the grain scale) and to have less of an effect on the residual strength of the material, this may be due to the fact that the kinematics of grains in the band are in any case very localised and intense, and that the material is flowing to some extent in a critical state. Using the analogy of flow, the secondary rotations could perhaps be considered as an effect similar to a boundary layer that retards the flow.

The kinematics in the residual state of Hostun sand show considerable secondary rotations caused by interlocking that in turn limit the rotations of grains. This is considered to be the principal reason for the increased residual strength of this material compared to Ottawa sand, which has grains of the same material. Shear banding in Ottawa sand causes less secondary rotations of grains outside the shear band: the material “flowing” in the shear band is slowed less by the surrounding material, which means that to have further shear in the band takes less energy than in Hostun sand.

Shear bands in Caicos ooids are even thinner than in Ottawa sand, with very limited secondary rotations. Based on these grain-scale observations Caicos ooids are expected to have an even lower residual strength: however, at both confining pressures studied, the residual strength is very close to that of Ottawa sand. This is most probably because of the different nature of Caicos’ carbonate grains, which have a higher inter-particle friction than the quartz grains of the other two materials; this could also explain the high peak of the Caicos material.

The polarisation of the grain rotation axes in the shear band, shown for Caicos and Hostun specimens at 100 kPa, reveals that the rotations in the increment studied are considerably more polarised for Hostun sand, but also relatively scattered as compared to the rotations in the Caicos specimen. However comparison is rendered more difficult by the fact that the increments studied are not of the same size. Over large increments, the expectation is that grains roll in the direction of the slope of the shear band, which can be seen in Hostun sand.

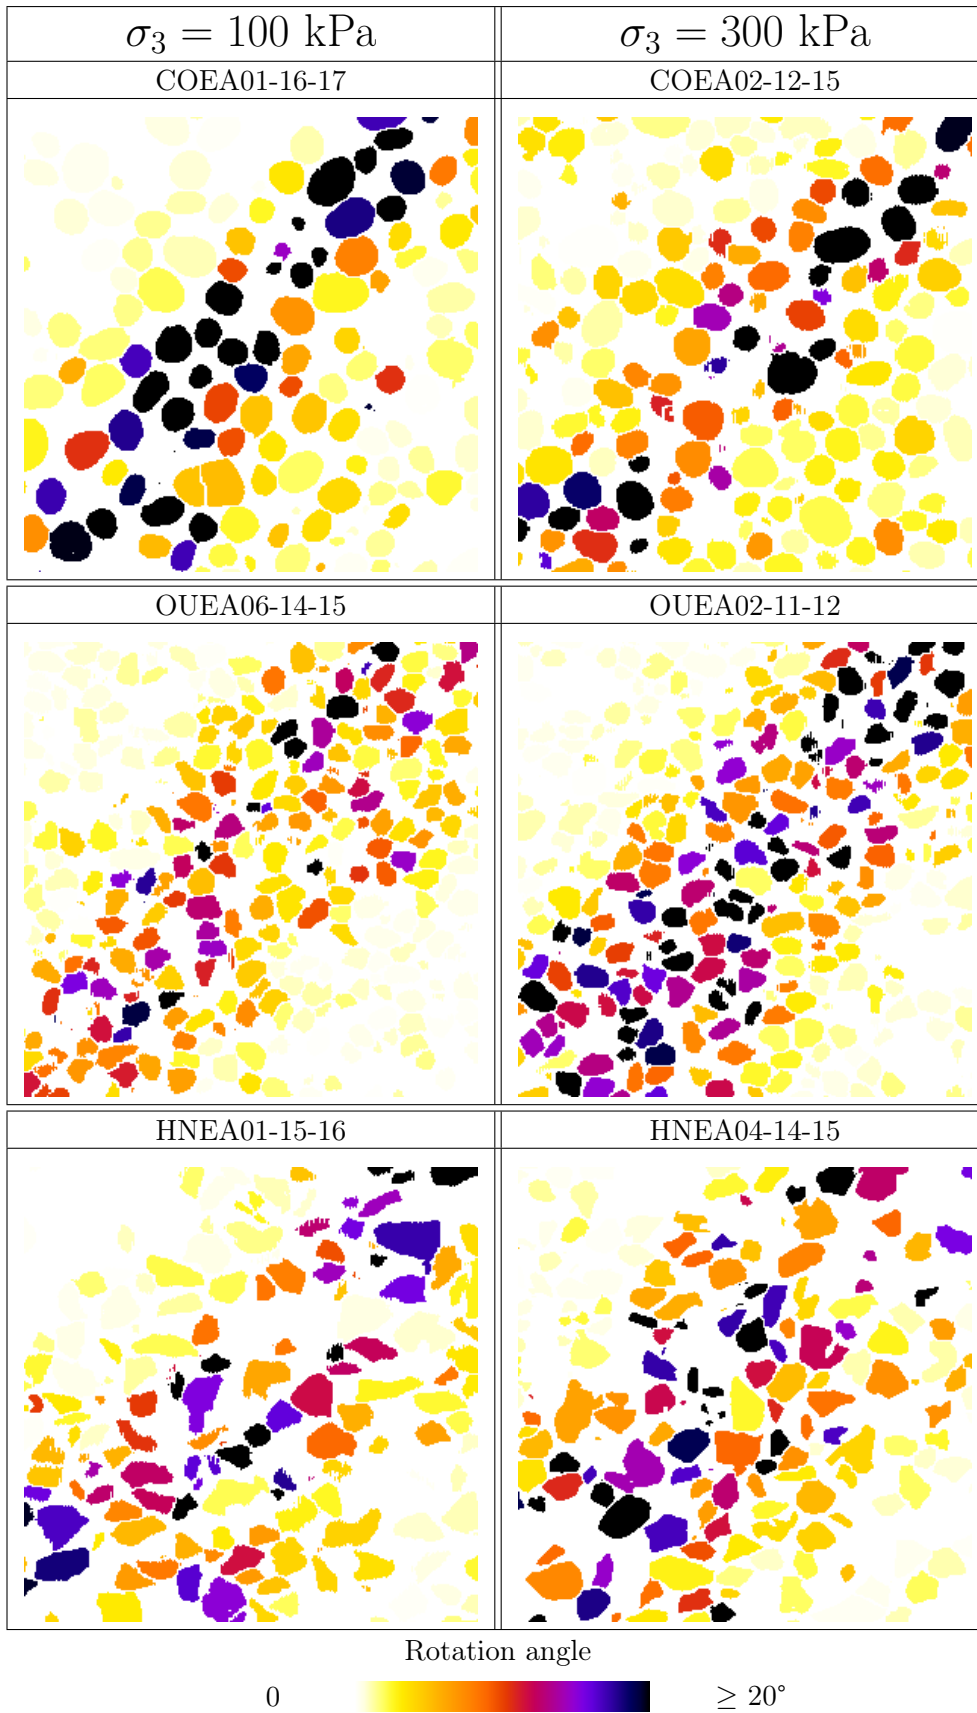


Figure 6.32: Comparison of grain rotations in the shear band

6.6.4 Development of shear bands

The section above has made a micro-mechanical characterisation of the mature shear bands that have developed in some of the specimens tested in this work, highlighting a potential link between the amount of secondary rotations outside the shear band with the residual strength of the material. This section endeavours to characterise the onset and progressive development of the shear bands.

Analysis of tests at 100 kPa confinement

The axial stress responses discussed at the beginning of this section show a remarkable trend in that the rounder the material, the “faster” the macroscopic stress response appears to be. For example, at 100 kPa confinement, COEA01 reaches peak stress at 3.6% axial shortening, and the residual stress state at 11%, but indications of the material being very close to the residual stress state come as early as 9% shortening. In contrast HNEA01 reaches its stress peak at 5% axial shortening and a residual stress state around 13.3% axial shortening.

The beginning of this Chapter has shown that in specimen COEA01 shear strain and grain rotations do not localise in the very first increment of the test but start to do so soon afterwards. In the residual state of the specimens analysed, the macroscopic stress response really reflects the material behaviour in the shear band. Between the beginning of the test (where no localisation is taking place) and the residual state (where the only active mechanism is the shear band, and a small thickness of grains around it), the shear band develops, and is likely to be responsible for the peak and strain softening observed.

The in-depth analysis of test COEA01 has revealed that porosity increases in the zone of the shear band, whereas there is little volume change on the outside of the specimen. This in turn has been related to coordination number of grains in the band, which have been shown to reduce (consistent with the increase of porosity) as the test advances. The fact that dilation is visible in the zone of the shear band is not surprising; the intense grain rotations and intense shear strain (differential displacements of grains) measured in the residual state requires enough geometrical freedom to be created for these mechanisms to occur.

Desrues *et al.* (1996) show that specimens of Hostun sand starting from different densities achieve a characteristic void ratio in the shear band that develop. These measurements have been repeated on specimen HNEA01 and are shown in Figure 6.33 where the porosity in the region of the shear band (red line) is plotted alongside the specimen’s global porosity (orange line). These results go some way to confirm the findings that a local limit void ratio for Hostun sand in the shear bands is 0.85 ± 0.02 (which corresponds to a porosity of $46\% \pm 0.6\%$) at 100 kPa confinement. As Figure 6.33 shows, specimen HNEA01 seem to be tending towards a porosity of about 48% in the shear band.

The conclusion to be drawn is that the porosity in the region of the shear band increases from very early on in the test, and tends towards a stable and well-defined value for the specimen. The final value of porosity for a given material is very likely to be related to the grain scale mechanisms identified in the previous section: for grains in the band to rotate and displace relative to each other, a certain amount of interlocking must be freed. A hypothesis could be that, until the critical void ratio or porosity required for the shear band to form is reached, the phenomenon of interlocking giving secondary rotations causes the material to have a stronger response than when the shear band is fully developed. Starting from the observation that increasing grain roundness seems to give increasing abruptness of the material’s response, the rate at which this increase of porosity can happen is expected to be higher for the more rounded grains.

Considering the other tests presented in Figure 6.33 shows some important are apparent: the rates at which porosity locally increases in the zone of the shear band are not the same between the tests, and the point at which porosity becomes stable in each specimen is also not the same. The values measured from the graph are shown in Table 6.4, showing that the Caicos specimen




Figure 6.33: Evolution of the mean porosity in the zone that the shear band occupies by the end of the test, as well as the global porosity for COEA01, HNEA01 and OUEA06




Figure 6.34: Macroscopic specimen responses for COEA01, HNEA01 and OUEA06 all tested at a confining pressure of 100 kPa. On the top plot, the stress ratio (R) is plotted against axial shortening, and on the bottom volumetric strain is plotted against axial shortening

is able to reach a critical porosity of around 47% porosity by 8% axial shortening by increasing

its porosity by 12.7%. As expected, Ottawa sand takes an intermediate place with a gain in the porosity in the zone of the shear band of 11.6% to reach a porosity of around 44%. This is obtained at around 10.6% axial shortening. Hostun sand is the material that gains the least porosity and reaches a stable porosity latest in the test.

More importantly, the maximum rate of increase of the porosity in the zone of the shear band reveals that for a given shortening, the Caicos ooids specimen increases in porosity more rapidly than the other two materials. Again, grain interlocking is seen to play an important role in the rate at which a localising granular material increases in porosity: as it has been shown in the section above, for a given axial shortening, the rounder grain have much less rotational frustration, and are therefore able to attain their “critical state” more rapidly.

Specimen	Max rate of porosity increase (increase of porosity per % axial shortening)	Approximate beginning of plateau (axial shortening %)	Absolute change in porosity over entire test (difference in porosity %)
HNEA01	1.2	11.7%	9.1%
OUEA06	1.5	10.6%	11.6%
COEA01	2.2	8.0%	12.7%

Table 6.4: Table showing some key measurements made on the evolution of porosity in the zone of the shear band in specimens tested at 100 kPa

Looking at the evolution of the coordination number for the grains in the zone of the shear band (Figure 6.35) the same trends are visible as for porosity: COEA01 is the test in which contacts are lost most rapidly in the zone of the shear band, allowing the specimen to locally reach an apparently stable value of coordination number. The other two specimens also show a behaviour that corresponds well to the observed evolution of local porosity discussed above. HNEA01 loses contacts in the zone of the shear band the least rapidly, therefore taking longest to achieve its stable state. As discussed earlier in the chapter, the measurement of coordination number appears to be more sensitive to changes in the macroscopic response of the specimen, however porosity still allows some interesting insights to be gained.

In both the porosity and coordination number maps, the Ottawa sand grains appear to have lower porosity and higher coordination numbers than the other two tests. This is likely not to be far from the truth, however, given the fact that the Ottawa grains are smaller than the other two materials tested, there are more partial volume voxels in the greyscale images, meaning that if these are selected by a very small change in the threshold used, a considerable number of contacts can be artificially identified.

In summary, at 100 kPa the roundness of the materials tested is strongly correlated both with the local rate of increase of porosity (and consequently loss of contacts) and also the extent to which local porosity increases over the initial state. When strain starts to localise, the medium’s macroscopic response is increasingly governed by the shear band occurring in the specimen (until the residual stress state, where the response is totally governed by the shear band). The more rounded materials tested appear to be able to reach a fully developed shear band “faster”, which explains why the peak stress is achieved sooner.

Figure 6.36 shows the deviatoric strain calculated in five key increments for these three tests:

- A: the first increment close to the maximum slope of increase of porosity in Figure 6.33
- B and C: two at the maximum slope of local porosity increase; one at the beginning and one at the end
- D: one increment when the gradient of porosity reduction starts to change

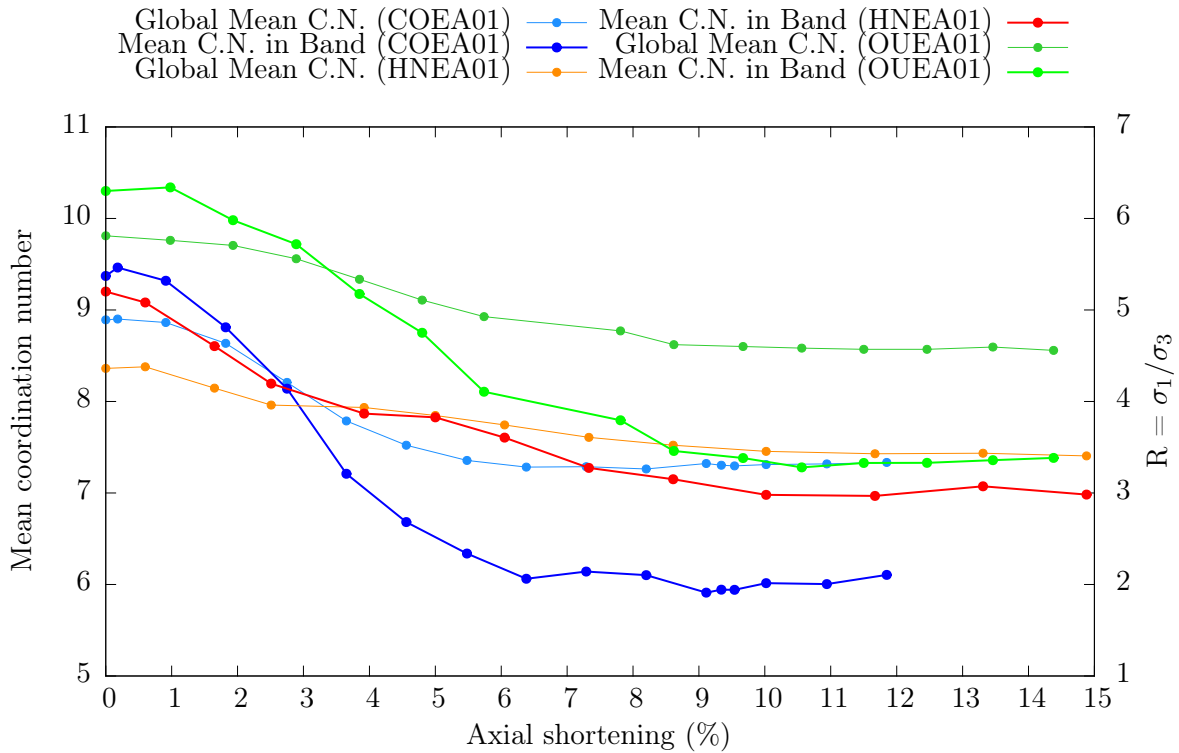


Figure 6.35: Evolution of the mean coordination number in the zone that the shear band occupies by the end of the test, as well as the overall mean coordination number for specimens COEA01, HNEA01 and OUEA06

- E: a last increment at the beginning of the plateau of local porosity increase

This figure shows that the beginning of the local change in porosity has a clear meaning: some very slight, diffuse features of localised strain can be seen for all specimens in increment A. Increments B and C, chosen from Figure 6.33, lie for every specimen on either side of the peak. For COEA01, in states B and C some considerable conjugate structures are visible, which are not very visible in specimen OUEA06. In increment C, HNEA01 shows some interesting features: two parallel structures delimiting the edge of the band, which can also be discerned in COEA01. What is noticeable in all specimens for increments B and C is that the diffuse zone of localisation visible in increment A has concentrated in space.

In increment D the shear band is in a much more recognisable form compared to each specimen’s residual state. By increment D, all the specimens have started the “second phase” of localised strain where strain is concentrating less in space and rather in value. By increment E the shear band is very close to fully mature: the second phase of strain concentration has crossed all of the specimens, and the increments all correspond to increments just at the end of strain softening.

Analysis of tests at 300 kPa confinement

At 300 kPa (Figure 6.37) the global and local evolution of porosity in the specimens qualitatively follows the same trends as at 100 kPa confinement: COEA02 has the greatest rate of increase, followed by OUEA02 and finally HNEA04, with a very low rate of increase of porosity. Table 6.5 summarises the same key measurements as Table 6.4 on these graphs. The values for HNEA04 are taken over the range 4% to 9% axial shortening to avoid the local change in gradient.

Table 6.5 (to be compared to Table 6.4) shows that the same qualitative differences can be seen between the three materials at 300 kPa as at 100 kPa, the values of specimen OUEA02




Figure 6.36: Slices through a field of deviatoric shear stresses computed from grain kinematics for some key increments from COEA01, OUEA06 and HNEA01




Figure 6.37: Evolution of the mean porosity in the zone that the shear band occupies by the end of the test, as well as the global porosity for COEA02, HNEA04 and OUEA02




Figure 6.38: Macroscopic specimen responses for COEA02, HNEA04 and OUEA02 all tested at a confining pressure of 100 kPa. On the top plot, the stress ratio (R) is plotted against axial shortening, and on the bottom volumetric strain is plotted against axial shortening

are at a considerably higher value of porosity than OUEA06. Comparison reveals that the

Specimen	Max rate of porosity increase (increase of porosity per % axial shortening)	Approximate beginning of plateau (axial shortening %)	Absolute change in porosity over entire test (difference in porosity %)
HNEA04	0.7	12.7%	8.0%
OUEA02	1.3	10.7%	11.0%
COEA02	1.6	10.0%	11.2%

Table 6.5: Table showing some key measurements made on the evolution of porosity in the zone of the shear band in specimens tested at 300 kPa

increase of porosity in the zone of the shear band is lower and slower for all the tests at 300 kPa confinement. Furthermore, the gains in porosity of Ottawa and Caicos specimens at 300 kPa are much closer together than the responses at 100 kPa. This is confirmed by the macroscopic responses (reproduced in Figure 6.38 for convenience), which clearly show a stress response that is very similar for the two specimens.

Conclusions

This section has shown that localised behaviour in the specimens starts early in the tests. The first indications of localised behaviour from porosity come early in the tests. Increments in which the local porosity starts to increase have been analysed at 100 kPa, showing that, in this early stage of a test, shear strain is already localising. The analysis of other increments, corresponding to key changes in the evolution of the porosity in the shear band, reveals that the period of maximum rate of increase of porosity contains the peak of the specimen's response and that the second phase of shear strain concentration starts when the rate of increase of local porosity reduces.

At higher confining pressure (300 kPa), the rate and amount of porosity increase is reduced for all specimens. The reduced rate of porosity increase is consistent with the "slower" specimen response seen at 300 kPa confining pressure. Observations at the micro-scale indicate that the local increase of porosity is initially related to the progressive concentration in space of the shear strain in the specimen. The maximum rate of increase of the specimen with respect to shortening coincides with the peak of the specimen's stress response and in some cases is associated with temporary, localised structures. When the rate of porosity increase starts to reduce, the shear strain in the specimen starts its second phase of concentration (no longer in space but in intensity), related to the end of the specimen's strain softening regime.

6.6.5 Initial Kinematics

This section will explore the kinematics of the initial increments of the specimens tested in order to establish whether the at-times imperfect geometry of the specimens has a strong effect on their response. The grain displacements, shear strain and grain rotation fields will be shown in the first increments of the six tests studied in this section. Since specimen COEA01 has a very small initial increment, for this specimen the increment COEA01-01-03 will be used.

The micro-scale measurements on the initial increments of six specimens studied in this work are presented in Figure 6.39. This figure shows that, even though the majority of these increments are well before peak (in general the end of the increment represents a state with a stress ratio of around half the peak), the fields of vertical displacement obtained by ID-Track already show some tilt. The specimens where the iso-values in the vertical displacement field are the most tilted are possibly the two Hostun sand specimens, HNEA01 and HNEA04 (the tilt of the bottom platen from horizontal was measured as 1.9° and 1.0° respectively). However, the deviatoric strain maps show that there is no localised shear evident in these specimens in their

Specimen	Vertical displacement	Deviatoric strain	Grain rotations
COEA01-01-03 $\sigma_3 = 100kPa$			
COEA02-01-02 $\sigma_3 = 300kPa$			
OUEA06-01-02 $\sigma_3 = 100kPa$			
OUEA02-01-02 $\sigma_3 = 300kPa$			
HNEA01-01-02 $\sigma_3 = 100kPa$			
HNEA04-01-02 $\sigma_3 = 300kPa$			




Figure 6.39: Micro-scale measurement made on the initial increments of six tests

initial increments. This can be confirmed by the absence of strong gradients in the displacement field. The rotation fields for these initial increments also show no noticeable activity.

Careful analysis of specimen COEA01's vertical displacement field in increment COEA01-01-03 reveals that, unlike the increment COEA01-01-02 analysed at the beginning of this chapter, the field is already tilted. This means that very early on in the test the iso-values of the displacement field can start to tilt. This in turn highlights the importance of the small increment (in terms of axial shortening) COEA01-01-02, which proves that the initial displacement field is homogeneous.

Figure 6.39 shows that despite the difficulty in the preparation of specimens, the localisation observed and characterised in these specimens is not caused by potentially poor initial conditions.

6.6.6 Conclusions of analysis of other tests

The analysis presented in Section 6.6 of the other tests performed in the experimental campaign reveals that the tools developed in this work are able to make measurements that enable the characterisation of grain-scale mechanisms in the different granular materials tested.

The questions posed (based on analysis of the macroscopic behaviour of these specimens) at the beginning of this section are recalled below:

- Can the stress response of specimens COEA01 and COEA04 be explained through micro-mechanical observations?
- Is the higher residual stress of the Hostun specimens explained by grain angularity?
- Is the relationship between grain roundness and the earliness of each specimen's peak explained by strain localisation occurring in the specimen?
- Can the lower and later response of specimens tested at 300 kPa confinement be explained by micro-mechanical observations?
- Is the different (carbonate) material of the Caicos grains only responsible for the high peak stress?

These questions have guided the analysis performed in this chapter, and have been answered in the discussion within this section; the answers obtained are summarised below.

The micro-mechanical comparison of specimens COEA01 and COEA04, tested under the same conditions, but giving a considerably different peak, first reveals that the residual state in these two samples is similar. Looking before the peak, no striking difference can be noticed until the peak. In the increment over the peak a mechanism of shear strain localisation was observed in COEA04, which took place much later in COEA01. The fact that the shear band was in a more advanced state of development in sample COEA04 means that the shear band brought the specimen to failure sooner, thus lowering the peak.

The systematically higher residual stress state in Hostun sand has been related to developments at the grain-scale and in particular, to "secondary rotations", which are hypothesised to be a result of rotational frustration of grains inside the shear band and thus indicate a material with more strength. Secondary rotations are seen to some extent in Ottawa sand, however in Caicos their presence is very limited. Based on the analysis presented in this chapter, the surprising residual state stress response is that of Caicos ooids, which when compared at a grain scale to Ottawa sand, is expected to have a lower residual stress state but instead systematically has a similar residual stress. This question should therefore be merged with the last question in the list: the fact that Caicos ooids are carbonate grains (as opposed to quartz) is probably the explanation for both the higher than expected peak stress, as well as the residual stress.

Analysis of the specimens' pre-residual micro-mechanical behaviour reveals a potential mechanism that explains the "speed" of their macroscopic response: the rate at which porosity increases (and that coordination number decreases) in the zone that contains the shear band at the end of

the test. Measurements made on the development of shear bands in Caicos ooids show that this material is able to locally increase in porosity almost twice as fast as Hostun sand. The analysis of Ottawa sand at different confining pressures reveals why OUEA02 behaves very similarly to COEA02 at 300 kPa confining pressure: sample OUEA02 increases in porosity at a similar rate to COEA02.

The change in specimen's macroscopic responses between 100 and 300 kPa confinement may come from two different mechanisms. The reduction of the stress response has been suggested to be the result of a slightly reduced amount of secondary rotations visible in the shear band. The retardation of the responses has been related to the observation that at higher confining pressure the zone of the shear band systematically increases in porosity more slowly.

Chapter 7

Example applications

This chapter aims to briefly introduce two different applications of the data coming from this doctoral work. The data are rich (particle kinematics of a specimen of sand through numerous increments have been calculated) and rather unique; this has caused a certain amount of interest which has matured into several scientific collaborations with various universities in France and abroad. Two illustrative collaborations are shown in this chapter, selected because of the relatively advanced state of the collaboration.

The first application of these data is detailed in Section 7.1 and describes an ongoing collaboration with Jose Andrade and Ivan Vlahinić in Caltech (U.S.A.). This is part of a large project, with the vision to combine discrete element models with the experiments made in this doctoral work.

Section 7.2, describes a collaboration with Antoinette Tordesillas, David M. Walker and Andrew Cramer in the University of Melbourne (Australia). This collaboration aims to make new measurements starting from grain kinematics and abstracting these data into kinematical space; once this is done a network is built and analysed using tools from graph theory. This differs from the analysis performed so far, which is always in 3D space.

7.1 Preparation for grain-based discrete elements

The collaboration with Jose Andrade and Ivan Vlahinić in Caltech (U.S.A.) is part of an ambitious project whose aim it is to make a new generation of Discrete Element Models (DEM), allowing a more profound investigation into micro-structural properties of granular media. Instead of imposing simplified shapes as is typically done in discrete element modelling, one objective is to take advantage of 3D measurements coming from the kinds of experiments carried out in this work, and to model much more realistic grain shapes.

The project hopes to reproduce *individual grains* in the specimens tested in this work by creating *avatars* of each grain from the images coming from tomography. The objective is to generate avatars of the grains which are then suitable for use in a discrete element simulation. In order for a grain to be suitable for a discrete element simulation, its surface should be described in a smooth way in order to avoid locking and numerical trouble in the discrete element simulation. However, a representation which is too smooth will inevitably lose information, particularly with angular sand grains like those found in Hostun sand. The representation of the surface of grains used in this project is a description by NURBS (Non-Uniform Rational B-Splines), which are chosen both to allow a correct description of a grain's surface to be made, as well as being suitable for subsequent computation: the surface of a grain being described by a function, it is possible and direct to determine the normal to the surface and its gradient at any point for example.

Before grains can be described by NURBS, their surfaces need to be accurately described starting from the 3D images from tomography. The methods developed in this doctoral work

identify grains as part of the solid phase, which in turn has been defined by applying a threshold to the greyscale images (which describe a field of x-ray attenuation) coming from tomography. The result of the application of a threshold is that the limit between solid and pore phases is strictly and rigidly defined – this is desirable insofar as it permits the solid phase to be subsequently separated into individual grains; however this strict separation means that the limit between solid and pore is jagged, since it is defined on cubic voxels. The jagged surface of a grain defined in a binary image is clearly visible in Figure 7.1 a), where the individual voxels are rendered as cubes in order to emphasise this effect.

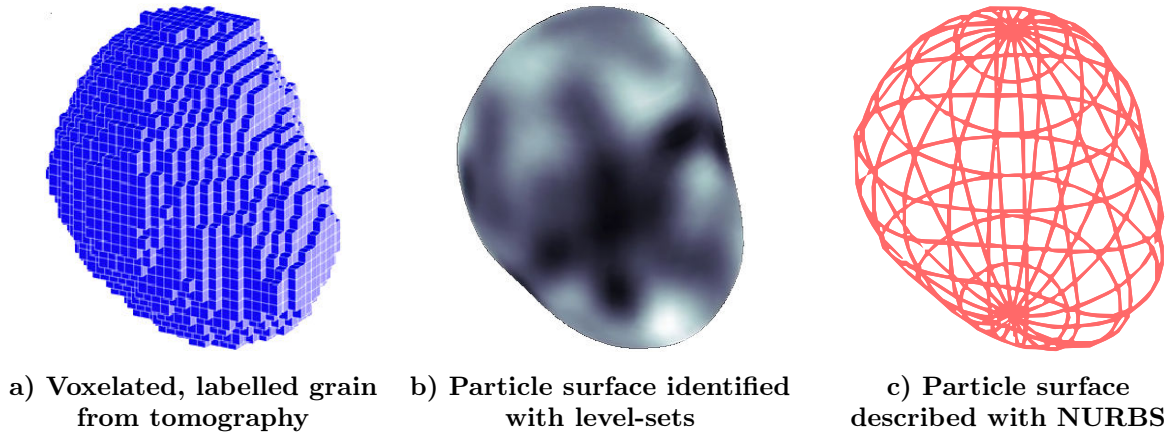


Figure 7.1: Figure showing the different representations of a grain – a) voxelised grain definition coming from the image processing steps described in Chapter 4. b) results of a level-sets procedure which smoothly defined the surface of the grain in question without being restricted to the resolution of the voxels of the image. c) smooth representation of the grain with a series of NURBS, allowing its surface to be captured efficiently (Jose Andrade, personal communication)

A paper has been submitted to *Granular Matter* (Vlahinić *et al.*, 2012) which describes a technique to better capture the surface of grains with the objective of a smoother description of the surface. The technique developed starts with the markers used to separate grains, but re-performs the separation using level-sets (see Osher and Fedkiw, 2003, for further details), guided by the greyscale images from tomography. Very briefly, level sets are boundaries that can be deformed based on image characteristics, such as the image value. Furthermore, some level-sets implementations allow pseudo-mechanical properties to be assigned to these boundaries, which can control their curvature or ability to bend. These boundaries are iteratively deformed in images until they become stable, and are not necessarily tied to the resolution of the image. In the implementation discussed in Vlahinić *et al.* (2012), the boundary expected to capture the outside of a grain is started from a marker inside the grain, and should grow until it stops on the outside of the grain. The partial volume effect on the outside of the grain is expected to help the placement of this boundary. This paper also develops some noise reduction filters for the greyscale image in order to guarantee a better convergence of these algorithms. These are illustrated in Figure 7.2 (parts “w/noise” and “w/o noise”).

The results of a correctly tuned level-sets boundary-finding/segmentation approach can be seen both in Figure 7.1 b) as well as Figure 7.2 (part “3D Level Set Contour”). Once a correct level-sets description is identified for a given grain, it can then be used as a scaffolding in order to construct a NURBS description of a grain. A NURBS description of a grain is shown in Figure 7.1 c) can then be used in the new DEM code being developed in this project. The NURBS description of the grain is its numerical avatar.

Once the technique is sufficiently fast to allow the avatars of a mechanically pertinent set of grains (*i.e.*, a representative elementary volume) to be calculated, the first objectives are to reproduce the micro-scale behaviour observed in the experiments by using DEM simulations on

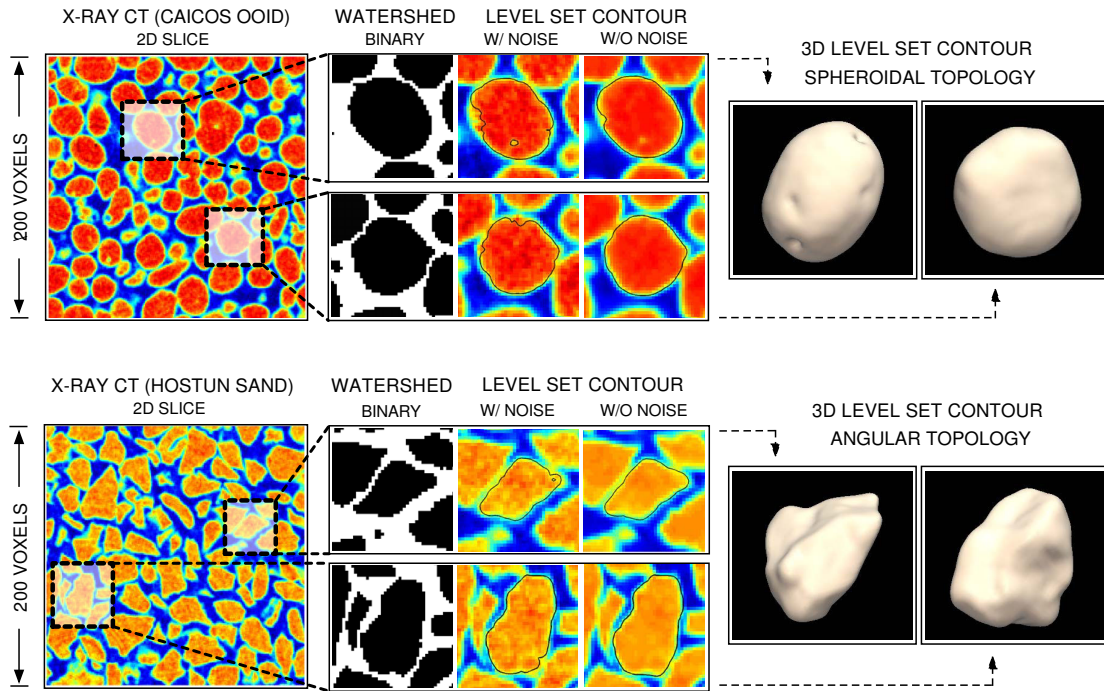


Figure 7.2: Figure taken from Vlahinić *et al.* (2012). Original caption: Characterization of rounded Caicoos ooids (top) and angular Hostun sand (bottom). (Left) A slice through a 16-bit 3DXRCT volume of 2003 voxels (approximately 1 mm^3). Representative grains varying from 193-253 voxels in volume were chosen at random from each specimen and segmented directly in 3D using the techniques described in this paper. (Middle) Grains shown at higher magnification. The proposed technique (Level Sets with prior image de-noising) captures the grain surfaces faithfully and with consistency. (Right) Grains rendered in Paraview (via a 3D Contour filter) without any topological modifications or enhancements

the grains' avatars. A 2D example coming from the collaborators (see Andrade *et al.*, 2012) of such a simulation is shown in Figure 7.3. This must be done by applying the same boundary conditions to the avatars of the grains in the DEM simulations. The same boundary conditions as the experiments could either be achieved by subjecting the avatars to the same stress path as the real grains were subjected to in the experiment, or (if the set of avatars is sufficiently large) to prescribe kinematics of grains on the limits of the cube of grains (with data coming from experiments), and compare the kinematics observed inside the box. The DEM simulation of the avatars needs to be correctly calibrated against the micro- and macro-scale response of the experimental measurement. This calibration must include confirmation of the microscopic response with other sets of grains from the same test, and then using the backup tests performed in this work. Once the DEM model is accurately reproducing the behaviour of the grains, the micro- and macro-scale response of an assembly of these grains can then be simulated on stress paths, and in geometries that are difficult or impossible to access experimentally. More importantly perhaps, the DEM simulations give access to the *forces* between grains, which opens a profound new line of investigation into the behaviour of real granular materials at the grain scale.

A DEM model simulating the response of “real” grains, which is able to capture the multi-scale response of a given granular material will first allow the many existing observations made with classical DEM to be checked and confirmed with a “real” material. Thereafter it will also allow simulations of the response of the granular material in configurations, stress states and geometries which may be experimentally difficult or impossible. This has the potential to become an extremely important tool for the investigation of the micro-mechanics of granular media.

Furthermore a finely tuned and realistic discrete element model is eminently suitable for use




Figure 7.3: Figure adapted from Andrade *et al.* (2012), showing a 2D DEM model capable of handling smooth grain shapes. Boundaries of the specimen are shown in orange, and grain perimeters are shown in blue

in the multi-scale simulations cited in the introduction. This modelling approach aims to replace classical constitutive relations in large-scale finite element problems with an explicitly simulated micro-scale model of soil behaviour.

7.2 Application of graph theory to analyse grain kinematics

This scientific collaboration with Antoinette Tordesillas and co-workers from the University of Melbourne focusses on the idea of using concepts coming from graph theory, and in particular recent developments coming from neuroscience, in order to look for abstract trends in the kinematical grain data produced in this doctoral work.

The fundamental hypothesis is that a large number of grains behave like a complex system. Mathematical tools from graph theory have been developed (with particular recent activity) for the study of complex systems (such as signals from neuroscience, or traffic flow) where trends and relationships are not immediately obvious. For example, the application of this kind of tool to data from neuroscience has revealed *functional* connections between parts of the brain which activate together – *i.e.*, which have a functional relationship between each other – but which are not close together in space. Furthermore, even when working in 3D space, graph theory can reveal some interesting trends in data.

In this collaboration, various measurements from graph theory are used in order to extract functional relationships from data of grain kinematics measured over various increments of test HNEA01 (which has been tested and analysed in this doctoral work). The principal findings have been written up and a journal paper has been accepted to appear in the Proceedings of the Royal Society: A (Tordesillas *et al.*, 2013). Some key results from this work will be illustrated in this section. Work is also ongoing to extract *structural* relationships from the HNEA01 data set, based on contacts between grains. Tools from graph theory in this case are expected to be able to easily pick out chains of contacting grains called “cycles”. The minimum number of cycles for stability in 2D systems is a 3-cycle (see recent work: Walker and Tordesillas, 2012); the number of 3-cycles is systematically seen to increase in shear bands.

The study of grain-based measurements in an abstract, functional space rather than in typical visualisation-in-3D space or just using specimen-wide averaging reveals some interesting trends. First of all, in order to be able to make this abstraction, a functional relationship between grains must be defined. This is done by the creation of a “*k-network*”, as illustrated in Figure 7.5. This is created by generating an *n*-dimensional kinematic graph for a given increment. In this graph, the kinematics (x, y and z displacement plus optionally rotation θ) of each grain is

used to plot a single point, or node, in this three- or four-dimensional kinematical space. This procedure removes structure from the information – the position of grains in the domain of the specimen is explicitly ignored. The nodes in kinematical space are then connected into a graph, by connecting, for each point, the k nearest neighbours. k is increased from 1 until all the graph becomes a “single connected graph”, which means that any node can get to all other nodes by following connections. The graph with the lowest k which is single-connected, is called the k -network for this increment. Figure 7.5 shows the creation of the k -network. This network can then be analysed, with a variety of metrics.

Highlights of two measurements made on the k -networks calculated on all the analysed increments of test HNEA01 (using measurement of displacement as well as a single angle of rotation) are shown in Sections 7.2.1 and 7.2.2.

Figure 7.4 recalls the macroscopic response of the specimen HNEA01.




Figure 7.4: Macroscopic response of specimen HNEA01

7.2.1 Community structures

Figure 7.5 (“community structure”) highlights the objective of the measurement of community structures: an algorithm attempts to separate the k -network into a certain number of different communities. The definition of a community in this work is a group of nodes with many intra-connections between them and few inter-connections to nodes on the outside of this group.

This is found by an iterative algorithm which is allowed to run and converge on all the k -networks representing all the analysed increments of test HNEA01.

Once communities are defined, these too can be analysed, and in this case the analysis focuses on two key features: first the size of the communities identified for each k -network are studied. Then, all the nodes belonging to a given community are projected back into the 3D space of the specimen and visualised; this allows groups of grains with similar kinematics to be identified in the domain of the specimen.

The size of a network can be expressed and measured in various ways, the simplest is probably the number of nodes making up a community. Figure 7.6 shows a different metric: “average shortest path length” – which is a measurement of size that is also sensitive to the layout of the nodes in a given community. This measurement is made by calculating for each node, the minimum number of links that must be followed to get to each other node. In this case it is measured for every pair of nodes in a given community and averaged, the averages for

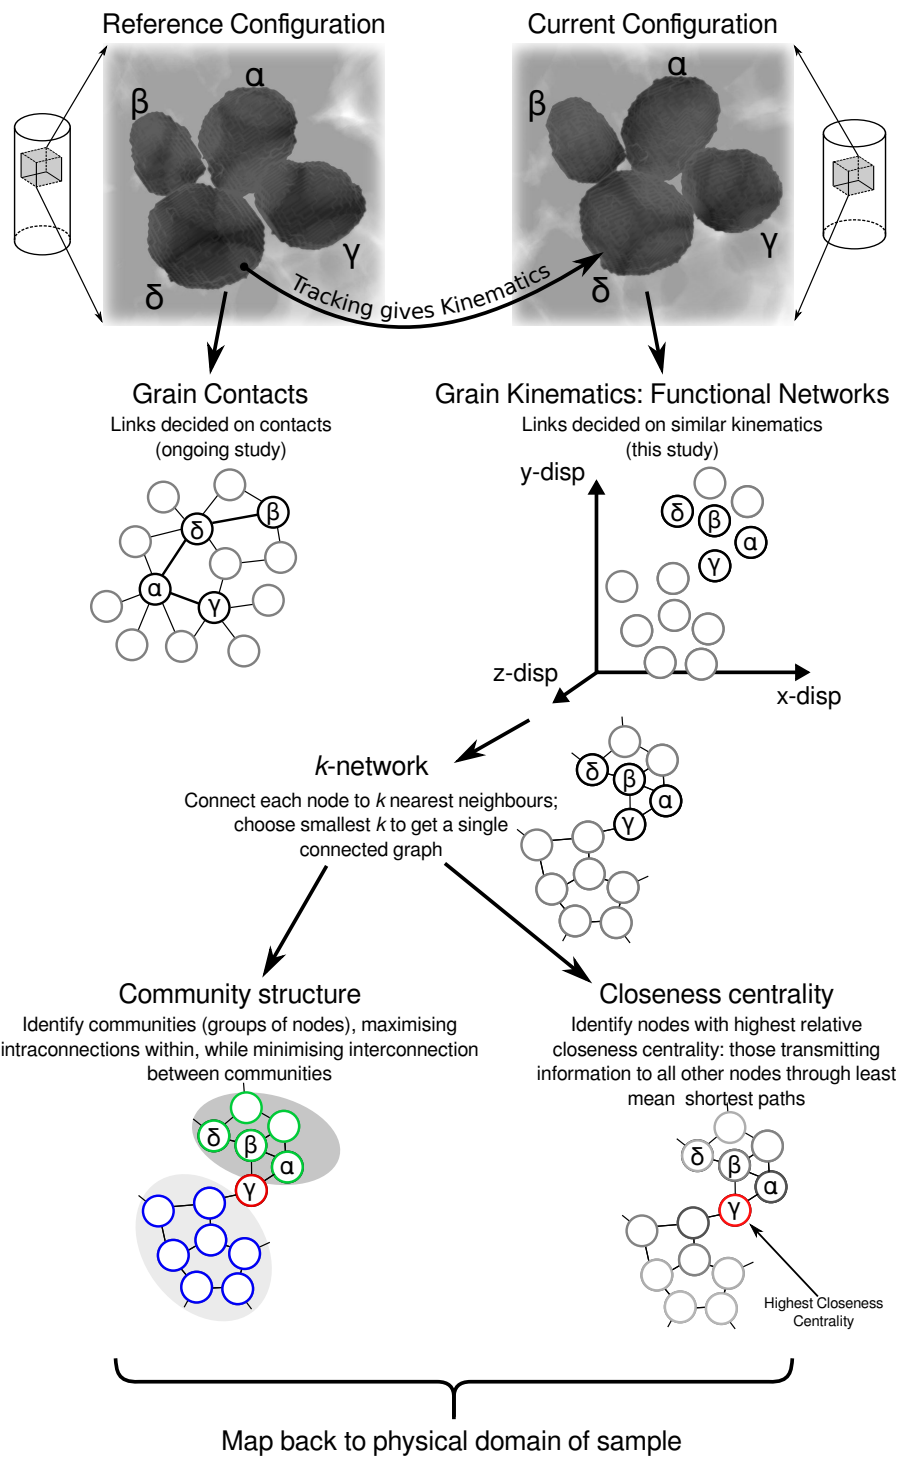


Figure 7.5: Figure taken from Tordesillas *et al.* (2013). Original caption: An illustration of our 3-phase strategy for characterisation of Hostun sand. Analysis starts from the physical domain of specimen (Phase 1). Individual grain kinematics are projected into kinematical state space (3D for grain displacement, 4D for grain displacement and rotation), and analysed as complex networks (Phase 2). A graph k -network is constructed: grains are represented by nodes, connected by links if corresponding grains exhibit similar kinematics, i.e. nearest k -neighbours in state space with respect to Euclidean distance. The network is characterised using various metrics, here centred around shortest paths. Information is then mapped back into physical domain of the specimen, trends are interpreted, and findings tested for robustness (Phase 3)

each community (shown in Figure 7.6) can then be studied. The median shortest path in each community (red lines in Figure 7.6) can be seen to increase until increment 10-11, the increment after the peak in the axial stress response of the test.




Figure 7.6: Box and whiskers plot of the distribution of shortest path lengths for all the communities identified in the k-network of the listed increments of HNEA01. The median is shown for each increment in red

The values observed in the median shortest path range from around 3 to 8 – which can be thought of as an internal length of the kinematic processes occurring in the specimen in each increment. The fact that this value increases around the peak of the specimen is interesting; it might point to the fact that communities of nodes (*i.e.*, grains) identified at the beginning and at the end of the test have very similar kinematics, and therefore very tight k-networks can be created – meaning that resulting communities will also be tightly knit, and therefore will not have long paths between nodes. The disorganisation of the kinematics as the specimen transitions from homogeneous deformation to highly localised deformation will cause a larger scatter of nodes in k-space, meaning that the k-network will be less compact, which would explain the increase in path lengths through the communities identified in this network.

In order to better understand the hypotheses made in the preceding paragraph, the information about nodes belonging to the same kinematical community are projected back into the 3D domain of the specimen. Figure 7.7 shows the communities identified in four different k-nets, projected back into the domain of the specimen in the configuration describing the beginning of the increment.

The first thing to be remarked is that nodes (*i.e.*, grains) belonging to the same community are in most cases also contiguous in space – but not in all cases. Furthermore, when the 3D space occupied by some communities is compared to the kinematics occurring in the specimen, it can be seen that the communities identified around the peak (increment 08–09 for example) have more diverse kinematics than the communities identified before and after peak. Before peak the kinematics are very similar all over the specimen, however, well after peak (say in the residual stress state) three main kinematic communities are expected: a community which is essentially stationary (the top of the specimen where displacement is imposed as zero), another community of grains which undergo a rigid body transformation (the bottom of the specimen) and a last one at the interface between these two, which represents a shear band. The image most on this right in Figure 7.7 shows this sort of behaviour.

This technique of abstraction of structural information from grain kinematics has allowed communities to be identified purely in the kinematics of particles. This has allowed the sizes of these communities to be measured confirming a length scale of around 8 connections at the peak of the specimen’s stress-strain response. Furthermore, the size of the community has been




Figure 7.7: Figure taken from Tordesillas *et al.* (2013). Original caption: Community structures ... mapped back into physical space (*specimen domain*) show that constituent grains form mesoscopic regions exhibiting similar kinematics. Spatial distribution of communities of the kinematical networks constructed using grain displacement and rotation for representative axial strain intervals of Hostun sand for 01-02, 02-04, 08-09 and 15-16, here shown as slices in xz plane. Grains with the same colour belong to the same community

related to the uniformity of the kinematics that a specimen is undergoing.

7.2.2 Closeness centrality

Another revealing measurement that can be made on the k-networks derived from grain kinematics is the calculation of the “closeness centrality” of each node. This is a measurement of the facility with which each node in a network can communicate with all the other nodes in the network. Nodes on the edge of a network will typically have to follow many connections in order to reach a node on the other side of the network. Nodes with many connections in the middle of the network are expected to be able to reach all the other nodes in the specimen with fewer connections.

The nodes identified with this technique as having high closeness centrality are expected to be the nodes which have kinematics (and not positions) which are in the middle of the principal kinematics of the specimen. In order to see which grains are identified with this procedure, the measurement of closeness centrality of each node is projected back into the domain of the specimen, and grains with a high closeness centrality are studied. Figure 7.8 a) highlights grains in red which have a closeness centrality in the upper 10% percentile with respect to the others. This is shown for two increments before the peak, one over the peak and one after the peak.

Figure 7.8 compares the results of the identification of high-closeness centrality nodes with results of Continuum DIC (TOMOWARP) showing a measured field of shear strain. In the middle two increments (02–04 and 08–09) both measurements appear to be sensitive to the shear band which is developing in the specimen, identifying the zone of the specimen in which the shear band develops by the end of the test in increment 15–16. Interestingly, in the increment 08–09 the maximum shear strain as measured by DIC appears to pick up some shear strain in the bottom right part of the specimen, whereas the closeness centrality measurement appears only to have identified nodes which are around the final position of the shear band.

The measurements in the first increment (01–02) appear to be relatively noisy, however the closeness centrality measurement appears to pick out the orientation of the shear band relatively accurately.

High closeness centrality grains in a k-network constructed from kinematics of a large set of grains is expected to identify the grains which are most able to communicate with other grains in that network. This measurement appears to identify, from an early stage in loading, grains which

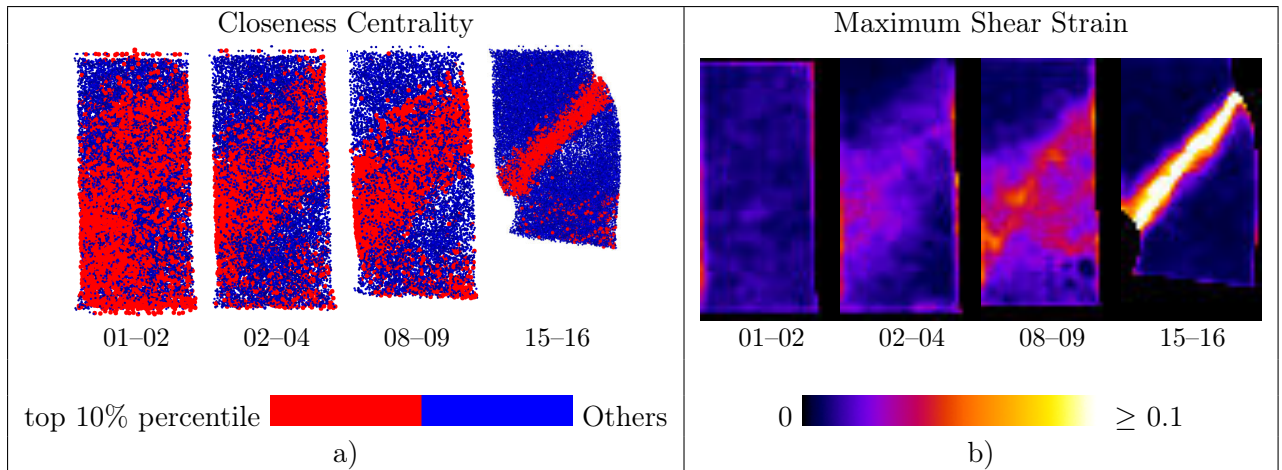


Figure 7.8: Figure adapted from Tordesillas *et al.* (2013). Original caption: (a) Centres of those sand grains corresponding to nodes of highest (upper 10% percentile) relative closeness centrality within the displacement-rotation kinematical networks, are mapped back into the physical domain of the Hostun sand specimen and shown as red dots for four strain intervals: from left are side-on views of 01-02, 02-04, 08-09 and 15-16 from the same perspective as those in Figure 2. The remaining grain centres are shown as blue dots. (b) Spatial distribution of maximum shear stress from grain-scale continuum strain measurements using digital image correlation, corresponding to the four strain intervals in (a). Notice the hint of a possible competing band for interval 08-09 in (b), just before the peak stress at 09. In contrast, the closeness centrality identifies only the final persistent shear band

have a particular significance – in the case of this data these are the grains which eventually end up in the shear band. The predictive ability of this measurement is certainly of great interest in order to uncover the reason for which it is able to predict the plane along which the specimen fails from the the early on in the test.

In order to check the power of the measurement of closeness centrality, the same exercise was performed on a discrete element simulation of spheres under plane strain. As before, kinematics of grains were calculated on small increments and k-nets constructed. Closeness centrality was calculated for each node in the k-network and then projected back into the physical domain of the specimen.

Figure 7.9 shows closeness centrality (with a normalised scale of closeness centrality, unlike the thresholded closeness centrality shown in Figure 7.8 a) calculated for each particle in a DEM simulation. Even right at the start of loading, a) shows that the closeness centrality measurement correctly predicts the position of the shear band. The confirmation of the predictive ability of the closeness centrality measurement from discrete elements (where initial boundary conditions are both easier to control as well as much more perfect than the experiments in this case) means this is indeed an interesting tool for the understanding of the behaviour of granular media.

The initial kinematics of the specimen appear to contain the final “fate” of the material, in the two examples developed in this work. If this can be systematically verified for very small loading increments, the next step is to try to understand what the initial kinematics depend on. Furthermore the robustness of this kind of measurement is also an interesting topic of study: the change in the angle of the predicted shear band in systems where principal stress rotation can be applied may reveal the mechanisms by which the failure mode can change in a granular material.

7.3 Conclusions

The two example applications of the data coming this doctoral work have been presented. These two particular examples have been selected because of the relatively advanced state of the work as




Figure 7.9: Figure adapted from Tordesillas *et al.* (2013). Original caption: Relative closeness centrality for the kinematical networks constructed using displacement and rotational information for six representative axial strain states, mapped back into the physical domain of the baseline specimen: (a) at start of loading, (b) midway through strain-hardening regime, (c) onset of force chain buckling and shear banding, (d) from peak shear stress to strain-softening regime, (e) transition from strain-softening to full development of the shear band, (f) a brief period of recovery in load-carrying capacity, evident in the rise in shear stress, after the shear band is fully developed. Note that the values of closeness centrality have been normalised to lie in the range $[0,1]$, i.e. the maximum value of closeness centrality is scaled to be one and the minimum is scaled to be zero with respect to the values at the considered axial strain and not across all strain states of the loading history. Increasing values of relative closeness centrality are depicted by particle colour ordered from cool to hot, i.e. from lowest (0) in dark blue to highest (1) in bright red

well as the promising results obtained; shown by the fact that in both cases journal publications have been submitted, and in one case accepted.

This chapter hopes to have shown the very different types of further work that are possible with the data acquired during this doctoral work; one approach aims to digitally reproduce the specimens tested in this work in order to reproduce the experiment with discrete elements in a numerical simulation. This means that the discrete element simulation can be compared at the grain scale to the experiments in order to best match reality. Once this can be done, the numerical model can be studied in order to look at the evolution of forces in the experiment. Furthermore the numerical model can then be used to investigate other situations which may be challenging or impossible experimentally.

The second application of these data is radically different, aiming to “data-mine” the rich data of grain kinematics abstracted from its positional information in order to draw out trends hidden by the “structural” relations in the data.

Both of these applications show promising results, and potential to further the understanding of behaviour of granular materials at the grain scale.

Chapter 8

Summary, Perspectives and Conclusions

This doctoral work has investigated micro-mechanisms of deformation in three different granular materials at different mean stresses through an experimental campaign. Access to the micro-scale has been enabled by x-ray tomography, which allows several 3D images of a specimen to be obtained as the specimen is deformed in triaxial compression. Techniques have been developed for measurement and extraction of grain-scale information from the 3D images coming from tomography. These tools allow, among other things, the displacement and rotation of all the grains that make up a small cylindrical specimen (more than 50,000 grains) to be followed through the full loading sequence (including, in general, 3D images of about 15 granular configurations for each test). This in turn has allowed the macroscopic response of these specimens to be related to the grain-scale mechanisms at play during their deformation. For example, the residual strength of each material have been linked to the amount of “secondary rotations” measured around the developing shear bands, which are strongly dependent on grain interlocking, and therefore grain angularity as well as the inter-particle friction. The point at which each sample reaches its peak stress has been related to its ability to locally increase in porosity, which again changes for each granular material.

This chapter concludes this thesis by first presenting a summary (Section 8.1) of the main points in the order that they have been presented. The chapter then continues onto the perspectives for the future coming as a result of this work (Section 8.2). This section is divided into the improvements that can be made to the tools developed; other questions that may be answered with these tools; further questions in which x-ray tomography may give considerable insight with the development of new tools and finally perspectives of work further in the future. The chapter closes with some final conclusions in Section 8.3.

8.1 Summary

An advanced experimental campaign characterising the triaxial compression behaviour of three different materials (Hostun sand, Ottawa sand, and Caicos ooids) at different confining pressures was undertaken. This experimental campaign was performed *in-situ* in the Laboratoire 3SR x-ray tomograph, *in-situ* meaning that specimens have been imaged in various configurations as they deform under triaxial compression. The specimens tested are much smaller than standard triaxial specimens (they measure approximately 22 mm height by 11 mm diameter). The zoom-effect of the conical x-ray beam is such that individual grains can be seen in the images coming from tomography. The preparation of small triaxial specimens has required an experimental protocol to be developed to ensure a consistent standard of specimen preparation. Specimens are created by depositing grains by air pluviation of the grains and are then tested dry.

A very simple analytical tool has been implemented to automatically optimise the geometrical settings required for the reconstruction of the 3D images of the specimens; the systematic

application of this tool to subsequent reconstructions of the same specimen reveals that there are geometrical changes in the acquisition system during a test. These geometrical changes were studied further by performing a test of the geometrical stability of the image acquisition system over several days. This revealed a strong correlation between the temperature of the x-ray cabin and differential displacements measured. Temperature sensitivity was studied and quantified in more depth with a specially-developed device, named TomoFrame. This device is a hollow rectangular frame with a steel sphere at the end of a short leg protruding from each corner of the frame towards the centre. The frame is designed such that thermal strains are compensated in all directions, so that the four spheres never move relative to each other. The TomoFrame can then be imaged by x-rays and a simple technique has been developed to follow the spheres in the radiographs. This allows the entire 3D rigid-body transformation from an initial state can be calculated and consequently allows the differential displacements in the acquisition system to be characterised. The TomoFrame can be installed such that the four spheres are projected into the corners of a radiograph, without disturbing a scan, which means that individual radiographs can be corrected back to an initial geometry. An example application shows a considerable improvement in the reconstructed image quality for long scans with small pixel sizes.

3D images acquired during the experimental campaign have been treated with relatively standard image-processing techniques to define the extents of the solid and pore phases. A porosity field can be calculated with relative ease (once the correct size of a representative elementary volume is established) in a 3D image describing the solid and pore phases. To define grains in the images acquired, the solid phase is split into individual grains using a marker-controlled watershed algorithm. This algorithm removes a surface of unit thickness between grains to isolate them, meaning that each grain can subsequently be labelled and properties such as its centre of mass, 3D volume, surface area and orientation can be calculated.

The surfaces removed between grains are considered as “contact voxels” and are also labelled using an algorithm developed in this work, giving each contact voxel between the same two grains the same label. Fitting a surface to the point cloud defined by all the contact voxels between two grains allows the contact orientations to be measured. Unfortunately the contact orientations obtained using this approach contain large artefacts due to the watershed algorithm used. A technique to improve the measurement of contact orientations is under development with Hugues Talbot and Clara Jacquet at ESIEE, Paris. This technique shows very promising results: a paper has been submitted to a conference on mathematical morphology (Jaquet *et al.*, 2013).

A grain-tracking technique named ID-Track has been developed in order to make the link between grains in two different configurations once grains in two volumes from an experimental sequence are labelled. The output of ID-Track allows the measurement of grain displacements from one imaged configuration to another. Grain displacements in turn can be used to calculate strain; of the several techniques possible one based on Bagi (1996) and implemented in YADE is used in this work. This technique involves the triangulation of grain centres with tetrahedra, on which a strain tensor is then calculated.

Grain rotations can be computed by ID-Track, by comparing grain orientations between configurations, however the accuracy of this approach is low, because grain orientations cannot be defined uniquely. An enhanced technique to measure grain rotations has been developed that combines ID-track (which is fast) with grain-based 3D digital image correlation (which is accurate). Grain-based 3D digital image correlation has been built from the ground up in this work, allowing the creation of a high-performance image interpolator that has been specifically developed for repeated image interpolations. The hybrid ID-Track and grain based image correlation approach shows substantial improvement to the comparison of grain orientations, and while being slower than a simple comparison of grain orientations, a considerable speedup is provided by

the fact that tracking results are used from ID-Track.

ID-Track, grain-based digital image correlation and other tools have been applied to the images coming from the experimental campaign, allowing measurements to be made on each test at the micro-scale. A test on Caicos ooids at 100 kPa confinement (labelled COEA01) has been studied in depth at the grain and contact scale and is subsequently compared to the other tests performed.

State-wise analysis of the specimens tested reveals that increases in the bulk volume of dense specimens during triaxial compression are related to a local increase in porosity inside a well-defined band. Careful scrutiny of the initial distribution of the porosity of some of the specimens does not reveal inhomogeneities that explain the location of concentration of shear in the specimens. Following the specimens throughout their deformation reveals that in the first increment of all the tests grain kinematics correspond closely to homogeneous deformation, with almost no localised strain or grain rotations. Thereafter, grain-scale analysis reveals that in the specimens tested, shear strain is organised in space around the position of the shear band (but not yet localised) well before the stress peak. As specimens are progressively deformed, the shear strain concentrates in space, tightening progressively until a final band width is achieved during the specimens' strain softening regime. At this point a second phase of concentration is observed: the *intensity* of shear strain concentrates, starting in one point of the band, and progressively spreading throughout the band by the time each specimen reaches its residual stress state. Some of the specimens show conjugate structures to the band, both in the strain fields and in the fields of grain rotations. Inspection of the rotation fields reveals that the conjugate structures visible around peak are lines of rotating grains. These "temporary bands" are de-activated as the main band becomes increasingly concentrated after peak.

The micro-mechanical comparison between the tests allows some differences in the specimens' macroscopic responses to be explained. For example, the residual strength of the specimens has been linked to the amount of "secondary rotations" of grains around the shear bands. At 100 kPa confining pressure, Hostun sand (which has the most angular grains) shows considerable secondary rotations, causing its residual strength to be the highest of all the specimens tested. Secondary rotations in Ottawa sand can be measured, but to a lesser extent than in Hostun sand, which explains its lower residual strength. In Caicos ooids specimens, the secondary rotations are very limited, however the residual strength is approximately the same as for Ottawa sand. This has been explained by the fact that the grains in this specimen are not quartz and are likely to have higher inter-particle friction, which would also explain the higher peak stress of the Caicos specimens. At 300 kPa confining pressure the same trends are observed, and the slightly lower residual stresses are tentatively linked to reduced secondary rotations.

Looking at the development of the shear bands at the grain-scale in the specimens tested reveals that the early apparition of the peak (with respect to the axial shortening imposed) and residual stresses is linked to grain angularity through the materials ability to increase in porosity. The most rounded materials are able to increase in porosity "rapidly", whereas the angular ones are not, and consequently have a delayed peak and residual state in terms of axial shortening. Studying the localisation process in terms of the rate of porosity increase shows that in all the specimens, strain localisation occurs well before the peak.

The rich and unique data coming from this work has provoked some international collaborations, two of which have been illustrated. The first collaboration with Jose Andrade in Caltech (U.S.A.) has an ambitious goal: the individual description of each grain from mechanical tests with *in-situ* x-ray scanning using a smooth representation of its surface by NURBS, thus creating a numerical "avatar" of each grain. The objective of this is to then insert each grain into a discrete element simulation, which is expected to be able to reproduce the specimen's response. Once this is possible, this tool will be used to look at the development of forces in a specimen, allowing

a deeper understanding of the grain-scale mechanisms observed. Furthermore, this avatar of the material can then be used to explore configurations and stress paths that are difficult or impossible to access experimentally. In another collaboration, with Antoinette Tordesillas (University of Melbourne, Australia), grain kinematics are explored using concepts and tools from complex networks. Grains are abstracted by considering only their kinematics (without the “structural” information of their positions). A network linking grains based on their kinematics is then created and analysed, showing that grains with high “closeness centrality” can be identified from early in the experiment as those that will participate in the final shear band. This result was confirmed on discrete element simulations, revealing that closeness centrality is able to make correct predictions of the final position of the shear band very early in a test.

8.2 Perspectives

The main accomplishment of this doctoral work is that a work flow and toolset has been developed, from specimen preparation to measurement of grain rotations, that allows the investigation of the processes of deformation of a granular assembly at the grain scale. This opens myriad possibilities for analysis, some of which have been explored (to different depths) in this doctoral work, as well as in two collaborations briefly described in Chapter 7. This section will describe some of the perspectives opened by this work. First the improvements that can be made to the measurement techniques will be outlined. This will be followed in Section 8.2.2 by a discussion of other outstanding questions in the mechanics of granular materials where the tools developed in this thesis can be used for investigation of phenomena at the micro-scale. The perspectives will then continue in Section 8.2.3 with further investigations that require the existing tools to be modified or new tools to be developed. The section concludes in 8.2.4 with a discussion of longer-term perspectives.

8.2.1 Improvement of tools for micro-mechanical characterisation

The tools developed in this thesis enable the experiments carried out to be analysed at the grain-scale. Some tools have room for improvement in terms of their error-rate and accuracy, which are discussed in this section.

The procedure for identifying grains in the 3D images from tomography currently involves the *binarisation* of the image into a map of the solid and pore phases in the specimen. This map is then split into individual grains by the use of a watershed algorithm. This step is crucial: once individual grains are identified in an image, they can be characterised by measuring, for example, their positions, volumes and orientations and the contacts between them. Furthermore, with two separated and labelled volumes, grains can be followed with ID-Track, which opens the way for grain-scale analysis of processes.

The watershed algorithm used in this work (part of the Visilog[®] toolchain) was selected principally for its speed: with the large amounts of data collected a fast segmentation is required. This however comes at a cost: the segmentation is not perfect on several points. The first point of great importance is that the separation of the solid phase into individual particles does not always respect the granular reality: over-segmentation can result in a real grain being split into several particles in the separated images and, conversely, under-segmentation can cause multiple grains to be represented by a single particle. In the images acquired in this work, a small amount of both of these errors has been noticed to occur. These errors mean that the characteristics calculated on the resulting particles do not reflect the characteristics of the grains they are supposed to represent. Furthermore, in cases where over- and under-segmentation do not occur in the same places between images, the particles cannot be tracked by ID-Track. The source of these problems principally resides in the way in which the algorithm identifies the starting points (or “markers”) inside each grain. Visilog[®] uses the standard morphological operators (locally merged maxima

of euclidean distance map) to identify markers inside particles. This approach occasionally fails to put the right number of markers in the right places, particularly with angular particles like those found in Hostun sand. In order to improve the separation of grains, an improved method for the identification of markers is required. However, this may be challenging due to the fact that both over and under-segmentation currently occur, which means that the criteria for the identification of markers cannot simply be adjusted by merging close markers together: a higher performance marker identification strategy must probably include some advanced morphological information. Another potentially interesting method for the improvement of grain separation is to take into account the fourth dimension of the data. If the hypothesis that grains do not break is made, when ID-Track is not able to track a particle this could indicate an incorrect separation in one or both configurations. Using past and/or future kinematics of the untracked grains, the incorrect separation could be identified, and locally re-performed. This idea has been named by Cino Viggiani as “Kinematics-Aided Segmentation”.

Another error that finds its source in the separation algorithm is the orientation of the contacts obtained between grains that have been separated. The separation algorithm used introduces very strong artefacts in the measured orientations. A different separation algorithm was noticed also to give strong (but different) artefacts in the measured contacts orientation. A technique that shows encouraging results is in the process of being developed. Once this is able to give reliable and accurate measurements of grain orientations the algorithm for the following of contact kinematics will be tested and used for further investigation.

The ID-Track technique itself can also be improved. This has been discussed in detail in Chapter 5; in Section 5.3.8 several possible improvements have been discussed, for example the improvement in the measurement of displacement by refining the measurement of grain displacements, as well a more intelligent definition of the search window to better take into account local grain displacements.

The measurement of 3D rotations of grains can also be improved. The main improvements to be made here are for the speed of the technique: the minimisation algorithm that guides the image correlation may well be improved by taking into account the history of the steps undertaken, in order to converge faster. Another approach to speed up this approach is the porting of the code to GPU, however this has not been attempted because of the high individual complexity of the operations made on the individual units of calculation (*i.e.*, the grains).

8.2.2 Further investigations with the developed tools

The tools presented and developed in this work can be used in their current form in order to answer questions other than those coming from the experimental campaign in this doctoral work. Using the same experimental equipment, for instance, these tools could be used to explore different stress paths, such as proportional loading, undrained conditions and triaxial extension. With different equipment other types of test are also attainable, such as oedometric compression.

Another interesting parameter space in which to expand the analysis is that of different grain shapes. A proposal recently funded by NASA (*“Physics-based modelling & characterization of regolith and its interaction with spacecraft: A paradigm shift for NASA’s missions”*), in which Laboratoire 3SR is involved, will use the tools developed in this work to make micro-mechanical measurements on granular media intended to simulate regoliths found on other planets. The long-term objective is to help build better models of regolith behaviour that will in turn help improve the design of spacecraft intended to explore planetary surfaces. A slice from a preliminary tomographic scan of a potential candidate for a Martian simulant is shown in Figure 8.1 a)

The micro-mechanics of other non-soil particles are also of interest. Triaxial tests on glass ballotini have been performed, but unfortunately rotations cannot be measured in these images. A small extract from a slice of a reconstructed 3D image of glass ballotini is shown in Figure 8.1 b). The creation of glass ballotini with a characteristic internal structure would allow rotations to be measured, which would contribute towards the ease-of-comparison with numerical results

obtained on spherical discrete elements.

Grain tracking is a technique of relevance for a scale of measurement in which individual particles can be distinguished but where there is not sufficient information within a single element in order to usefully measure its evolution. At a resolution of observation where particles cannot be distinguished, or in situations in which there is sufficient information inside a given particle, continuum DIC is the most relevant tool to use to follow the material. Particle tracking is therefore useful at a certain scale of observation. A pixel, however, can correspond to very different sizes depending on the multitude of imaging systems that exist, and on the ability to zoom in systems like the x-ray scanner in Laboratoire 3SR. This means that particle tracking is relevant in a range of different physical scales; Figure 8.1 c) shows a slice from a tomography on a non-soil granular material (lentils) which has much larger grain size than the grains studied in this work; ID-Track has been applied to these images with success. This image was acquired on the Laboratoire 3SR tomograph, taking advantage of the ability to zoom. With a sufficiently capable x-ray system, particle tracking could also be applied to smaller specimens of smaller grains; for example, with recent developments in x-ray sources, testing on silts starts to become a possibility.

Other examples of non-soil applications for these techniques come from an ongoing collaboration for the study of mechanical compression of industrial products (detergents), a specimen of which is shown in Figure 8.1 d).




Figure 8.1: Images showing some other granular materials suitable for particle tracking. Image d) Courtesy of Jin Ooi, University of Edinburgh

Another interesting area of study would be to repeat the tests performed in this work, but using soils with a wider grain size distribution (while remaining in the range treatable by tomography), to investigate the relationships between the grading used and the development of shear banding for example. The surprising effect of grading and grain size on the development of shear bands reported in Viggiani *et al.* (2001) may well be explained by grain-scale processes.

8.2.3 Investigation of micro-scale mechanisms requiring further developments

Hopefully this thesis has shown to some extent the power of x-ray tomography for the investigation of the deformation for granular materials. The tools developed herein are suited to the following of permanent particles being macroscopically deformed. More generally, x-ray tomography can also shed light on different processes in geomaterials, a few examples of ongoing work in Laboratoire 3SR are given in this section.

An extension of the campaign performed in this work is underway to study the behaviour of sand at higher confining pressures; Figure 8.2 shows a zoom into vertical slices taken from a specimen of Ottawa sand at 7000 kPa confining pressure. The image on the left shows the specimen under isotropic compression, and the image on the right shows the same zone in the

specimen after 15% axial shortening: a compactive shear band (*i.e.*, region of crushed grains) crossing from top left to bottom right is visible.




Figure 8.2: Zooms of vertical slices taken from a triaxial compression test on Ottawa sand

It is of obvious interest to identify grain-scale kinematics in this sort of experiment, however since individual grains are no longer permanent features, some major modifications will have to be made to the tools developed in this work in order to take this into account.

Another series of experiments run in the Laboratoire 3SR tomograph consisted in the triaxial compression of bio-cemented Ottawa sand specimens (see: Tagliaferri *et al.*, 2011). The quartz Ottawa sand grains were cemented by calcite (which has a higher density), which means that the cement is visible and (more importantly) distinguishable from the sand grains in the 3D images acquired. Figure 8.3 shows a zoom into a horizontal slice of a bio-cemented specimen, where calcite cement (lightest grey) is clearly visible between grains. With careful thresholding, the *grain* phase might be identified in these images, so that they can be used with ID-Track. The cement, and, more importantly, the evolution of the cement with imposed shear is of great interest since it has a considerable mechanical effect on the response of the specimens tested. To be able to take the evolution of cement into account, it first needs to be identified reliably by facing some of the problems associated with image *trinarisation*. Once done, the contact kinematics algorithm developed in this work could be adapted to follow the cement between particles.




Figure 8.3: A zoom from a horizontal slice of a bio-cemented Ottawa sand specimen

The micro-scale investigation of water repartition in partially saturated sand is also a subject of ongoing work. This is related to the cemented granular materials described above insofar as

there are three distinct phases in the material being imaged: solid, liquid and void. In this work, the wetting curve of Hostun sand has been explored with a specially developed cell that allows control of suction inside a cylindrical 10×10 mm specimen. In the masters work of Riedel (2011) and Khaddour (2012) Hostun sand specimens were scanned at high levels of zoom in different states. The subsequent trinarisation of the images coming from tomography allowed the extent of the solid, liquid and void phases to be identified. Figure 8.4 shows a result from several steps of image treatment (see: Kaddhour *et al.*, 2013).




Figure 8.4: A trinarised horizontal slice of a Hostun sand specimen in a state of partial saturation, with the void phase in black, liquid phase in blue and grain phase in brown

A recent ambitious project has developed a triaxial cell like the one used in this work which adds the ability to control suction inside the specimen. Images from this project can also be studied at the grain scale, and like the case of the cemented sand discussed above, new developments will allow the following of individual grains as well as the water phase.

8.2.4 Future work

As cited in the introduction to this manuscript, Calvetti *et al.* (1997) clearly lay out the roadmap for experimental micro-mechanics. Future work in this field can be situated inside the framework that they define:

“The full description of micromechanical behaviour may be divided into three points:

- 1. Description of the structure, that is to say, position of grains and contacts between them.*
- 2. Description of the kinematics evolution: displacements, rotations, evolution of contacts.*
- 3. Description of intergranular forces.”*

With the development underway for the improvement of the characterisation of contacts, the techniques developed in this work cover points one and two. The third point is a considerable challenge in 2D and is of significant technical difficulty in 3D, however some in-roads have been made.

The masters work of Wiebicke (2012) is a 2D experimental verification of the “Granular Element Method” developed in Andrade and Avila (2012). The granular element method is able to give a contact force at each contact given knowledge of the boundary conditions, average stress tensor inside each grain and position of contacts between grains. In this work, 2D continuum digital image correlation was used to analyse images of 11 rubber “grains”, with dimensions in the order of 10 cm, in plane stress compression. Stress fields inside each grain were obtained by the using the strain fields coming from DIC combined with a simple constitutive relation, allowing GEM to be applied with encouraging results, giving force measurement at the contact points..

Some previous work with 3D x-ray tomography such as Saadatfar *et al.* (2012) has made first steps towards the mapping of the stress and contact forces based on measuring the strain in individual grains. The main limitation of this kind of approach, for the moment, is that resolution in the measurement of the strain of a single grain must be traded off against the number of grains; which currently means having specimens which are composed of few particles (around two thousand particles in the cited work). In 3D recent experimental work on sand grains, x-ray and neutron *diffraction* techniques were used to measure grain strains, again with encouraging, but very preliminary results (see Hall *et al.*, 2011).

An indirect method for obtaining forces has been described above in the collaborative work with Caltech (U.S.A.). This approach aims to reproduce, grain-by-grain, a specimen from its 3D image by creating numerical avatars of all the grains, which can then be used in a specially-developed discrete element method to infer the inter-granular forces.

8.3 Conclusions

This thesis has presented an experimental characterisation of the micro-mechanisms associated with deformation in three different granular materials, making measurements on small triaxial specimens whose micro-structure has been imaged in 3D by x-ray tomography in various configurations.

Techniques have been developed that allow measurement of each grain's displacement and rotation which can be related to the macroscopic stress and volumetric response of the specimens tested.

The main conclusion from this work is that grain-scale observations are possible on granular systems and that they are able to reveal many important aspects of what deformation means at the grain scale. This follows to a certain extent in Terzaghi's line of reasoning (fully cited in the introduction) in "*starting again from the elementary fact that sand consists of individual grains*". For example, the observations made in this work show irrevocably that shear banding in sands is a grain-scale phenomenon.

As said more recently by another important character in the geomechanics world, "*While we are happily ignoring particles, they will at times come back to haunt us. This happens when deformation is localized in regions so small that the detail of the soil's particular structure cannot safely be ignored. Failure is the perfect example of this*". Advanced modelling of strain localisation, and therefore failure, requires this small-scale to be taken into account. A promising line of developments are multi-scale models, in which different scales are explicitly simulated. For example, macro-scale finite element simulations with a separately modelled micro-scale represented by discrete elements (DEM-FEM models, *e.g.*, Nitka *et al.*, 2011) are particularly promising. For these models to succeed, the discrete element simulations they are based on must be set up such that they convincingly reproduce reality, which in turn requires systematic observations of the evolution of the micro-structure of real granular materials as they deforms under load. Hopefully the tools and results of this thesis are an important step towards this goal.

Chapter 9

Résumé en Français

9.1 Introduction

Cette thèse de doctorat se situe dans le domaine de la micro-(géo)mécanique expérimentale. Une campagne expérimentale d’essais de compression triaxiale à petite échelle a été réalisée sur des échantillons de sable, variant en forme de grain et de la pression de confinement. La nouveauté important que ce travail présente est que chaque essai triaxial a été effectuée à l’intérieur d’un tomographe à rayons-x (celui du Laboratoire 3SR, Grenoble), permettant l’acquisition d’un nombre de radiographies en faisant tourner tout le système de chargement, échantillon inclus, atour d’une axe de rotation verticale. Ce jeu de radiographies, acquises d’une façon non-destructive pendant 2h, permet de obtenir une images 3D de la microstructure de l’échantillon. Puisque les radiographies sont non-destructives, ces acquisitions peuvent être enchaînés pendant un essai, permettant donc un suivi *in-situ* (*i.e.*, des scans pendant le chargement, ce qui est rendu possible grâce à une cellule triaxiale spécialement conçue sans barres de retour de force – voir la Section 3.4).

Les échantillons testés ont été miniaturisés: ils mesurent 22 mm en hauteur et 11 mm en diamètre de sorte qu’ils peuvent être entièrement balayés par un faisceau conique de rayons-x à un niveau de grossissement tel que *chaque grain* qui compose l’échantillon peut être identifié dans les images *de champ* 3D provenant de la tomographie à rayons X. La représentativité mécanique de ces petits échantillons est traité dans la Section 3.7: leur réponse mécanique dans ces conditions de chargement se avère d’être bien représentative d’essais réalisés sur des échantillons classiques. Dans les essais de taille standard (et en particulier dans le cas des matériaux denses) la localisation se produit avec des bandes de cisaillement, comme il le fait dans ces essais à petite échelle. Pour l’étude de la localisation de la déformation dans les milieux granulaires, ces essais à petite échelle ont l’avantage considérable de permettre l’étude de l’apparition de bandes de cisaillement à *l’échelle du grain*.

Afin de reconstruire des images tridimensionnelles à partir des radiographies acquises pendant un scan pendant lequel l’échantillon est tourné à travers 360°, l’axe de rotation doit être connue avec une bonne précision. Puisque l’axe de rotation n’est pas parfaitement fixe pendant un scan (surtout pour des scans longs et à un fort zoom), un nouveau dispositif physique de correction (nommé le TomoFrame[®]) a été développé et breveté le 8th April 2014 (voir Section 2.3). Ce dispositif permet de mesurer les déplacements “parasites” de l’objet sous scan en 3D, et ensuite de les corriger. De plus, puisque sur le tomographe du Laboratoire 3SR, les déplacements parasites sont fortement corrélés avec la température, le positionnement des point de repère du TomoFrame[®] sont immobiles pour des changements de température.

Les trois matériaux testés sont, en ordre croissant de rondeur, sable de Hostun HN31 ($D_{50} = 338\mu\text{m}$), le sable d’Ottawa 50/70 ($D_{50} = 310\mu\text{m}$) et les “Caicos ooids” ($D_{50} = 420\mu\text{m}$) – comme expliqué dans la Section 3.2. Le principal de l’analyse mécanique dans ce travail de thèse porte sur une campagne de 13 essais de compression triaxiale sur des empilements granulaires denses

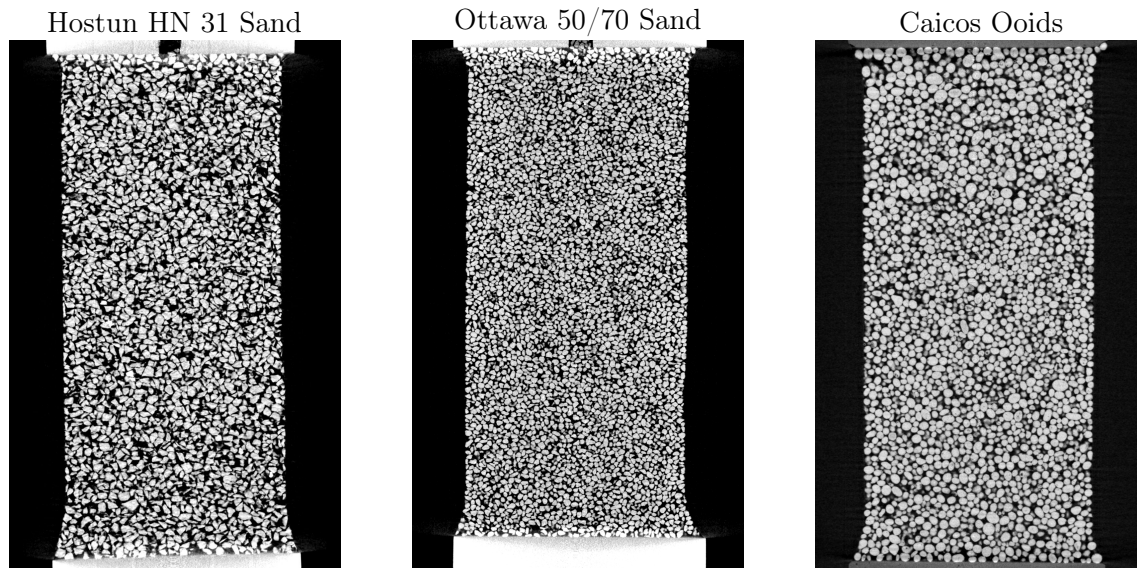


Figure 9.1: Coupes verticales extraites des images tridimensionnelles de trois échantillons (des trois types de grains étudiés) qui ont été déformés sous 100 kPa de confinement, sous 100 kPa de contrainte isotrope (*i.e.*, avant le début du chargement déviatoire, dans un état qui peut être considéré comme initial). Rappel: hauteur initiale est d'environ 22mm

pour chacun de ces matériaux. Ceci est obtenu par une pluviation dans l'air (voir Section 3.3). La campagne d'essais balaye deux paramètres: les différents types des grains, et la pression de confinement sous laquelle chaque essai de compression triaxial se déroule (pour chaque matériau il y a au moins un essai à 100kPa et un autre à 300/400 kPa). Les conditions initiales et paramètres de test pour chaque essai est détaillé dans la Section 3.6.

Pour chaque essai triaxial, environ 15 images 3D ont été acquises, toujours avec une première image dans un état de contrainte isotrope, et le reste des scans distribués pendant la compression déviatoire. Ce travail de thèse s'est concentré sur la réalisation de cette campagne expérimentale, et de l'analyse fine des données acquises pendant les essais: non seulement les mesures de contrainte et déformation macroscopiques, mais aussi le traitement des images issues de la tomographie afin de faire des mesures à l'échelle du grain, et mettre en relation l'évolution des ces micro-quantités avec les mesures macroscopiques.

9.2 Analyse d'images

Les images issues de la tomographie à rayons-x sont des images en niveau de gris (telles que celles présentées dans la Figure 9.1. Ces niveaux de gris représentent (passant pas quelques hypothèse assez fortes sur le faisceau de rayons-x) l'atténuation aux rayons-x dans un volume élémentaire (nommé voxel – l'équivalent 3D et donc volumique du pixel) reconstruit des nombreuses radiographies, ou projections de l'échantillon tournant dans un état donnée. Il y a une relation très forte, pour une énergie de rayons-x donnée entre l'atténuation de la matière et son numéro atomique, Z . Il y a aussi un lien entre Z et la densité – donc les cartes en niveau de gris peuvent être, plus ou moins, lues comme des cartes de densité.

A partir de ce genre d'images, et passant par une calibration des niveaux de gris de la phase solide et la phase vide, des quantités telles que la porosité peuvent être définis soit à l'échelle de l'échantillon – une porosité globale, ou à une échelle "mesoscopique", localement dans l'échantillon. Cette dernière définition permet, par exemple, d'évaluer l'homogénéité initiale de l'échantillon. La définition de la taille du sous-volume sur lequel faire cette mesure de porosité locale n'est pas facile à calculer directement – il y a un équilibre à établir entre une taille trop grande qui n'est pas sensible aux variations locales, et une taille trop petite qui est trop (et qui

introduit donc du bruit) sensible aux grains. Ceci est traité dans la Section 4.3.1, et la Figure 9.2 illustre bien le concept: des sous-volumes de la taille d'un seul voxel seront soit dans un grain, soit dans un vide (porosité 0% ou 100%). En croissant la taille du sous volume, la mesure de la porosité du sous-volume a la tendance de converger vers la porosité globale de l'échantillon.




Figure 9.2: Evolution de la mesure de la porosité fait dans un sous-volume cubique la taille duquel varie entre un voxel et un cube de $100 \times 100 \times 100$ voxels. Ceci est illustré pour l'échantillon HNEA01, dans l'état de 01 pour dix nœuds différents placés sur l'axe de l'échantillon et avec 200 px espacement vertical. Comme le sous-volume de calcul augmente, la porosité mesurée peut être perçue comme tendante vers une valeur représentative de la porosité globale de l'échantillon

Ce résumé va se concentrer sur les résultats *discrets* obtenu pendant cette thèse – c'est-à-dire les résultats attachés aux grains même qui constituent les matériaux granulaires qui sont le sujet des ces études mécaniques. Pour définir strictement la phase solide dans les images à niveaux de gris issues de la tomographie, un seuil est appliqué aux images – à des niveaux de gris inférieur au seuil le pixel est considéré étant partie de la phase vide, sinon le pixel est assigné à la phase solide. Le seuil est calibré pour donner un volume solide égal à celui qui est déduit du pesage des grains après-essai. Le résultat de l'application du seuil est une image dite "binaire" (c'est-à-dire avec seulement deux niveau de gris – 0 pour le vide et 1 pour le solide). Dans ce genre d'image les deux phases – solide et vide – sont les deux intimement inter-connectées.

Afin de définir les grains individuels dans l'image, la phase solide doit être séparée en objets (*i.e.*, particules/grains) discrètes. Ceci est obtenu en utilisant une implémentation commerciale du *watershed* algorithm (algorithme de ligne de partages des eaux en français), qui opère sur la phase solide. L'implémentation utilisée fait partie de Visilog[®], et cette implémentation ne fait partie des plus récentes qui propose des lignes de partage *interpixel*, ceci veut dire que les différents objets identifiés sont "manuellement" séparés dans l'image, en enlevant une épaisseur d'un voxel entre chaque objet – opération légèrement intrusive puisqu'elle enlève de la masse qui a été calibré attentivement. Partant d'une image de la phase solide séparée, chaque objet peut être facilement numéroté, ce qui donne une image "labélisée". L'enchaînement de ces différentes étapes d'analyse d'images sont résumées dans la Figure 9.3. Puisque l'algorithme sépare les grains en enlevant des voxels, là où des voxels ont été effectivement enlevés entre deux objets, ceci est considéré comme indication de la présence d'un contact, et ceci est enregistré pour du traitement subséquent (voir Section 4.1.4).

Étant chaque grain numéroté individuellement, c'est trivial d'aller chercher l'ensemble des





Figure 9.3: Illustration très synthétique de la méthode pour identifier les grains à partir d’une image à niveaux de gris: Binarisation (see Section 4.1.1), Separation (see Section 4.1.2) and Labelling (see Section 4.1.3)

voxels solides dont il est composée, et en calculer par exemple le centre de masse (calcul qui donne une précision bien en dessous de la taille du voxel). Les différentes propriétés intéressantes qui peuvent être calculées directement sur un ensemble de voxels (par exemple l’orientation des axes principales de l’objet) qui représentent un grain sont données et caractérisées dans la Section 4.2.

Fort des nombreux outils pour discrétiser et ensuite caractériser une image tomographique d’un milieu granulaire, dans le déroulement de cette thèse la réalisation que la numérotation des grains n’est jamais régulière est arrivée rapidement. Le réarrangement des grains et des petites incertitudes des images fait que les grains n’ont jamais les mêmes identifiants entre deux images. Pour résoudre ce problème, un des gros développements dans cette thèse est l’outil ou l’algorithme “ID-Track” (Publié dans Andò *et al.* (2012a)). Cet algorithme fait un suivi de grains entre états imagés, et reconnaît les grains par une hypothèse de peu de changement de volume de l’objet dans un premier passage, pondéré par une déviation par rapport au déplacement mesuré des voisins dans un deuxième passage. Avec cette mise-en-relation des différents identifiants des grains entre images, le suivi du centre de masse d’un grain donné entre images donnera son vecteur 3D de déplacement par rapport à l’origine de l’image.

RTEA01 Reference Configuration



Caicos ooid

Play-Doh

Ottawa 50-70 sand

Play-Doh

Hostun HN31 sand

RTEA01 Tilted by 4.7°




Figure 9.4: Des coupes verticales de l’essai calibrée de rotation

Les rotations des grains peuvent être calculées à partir du tableau qui résulte du suivi des grains. Pour pouvoir caractériser la performance des mesures de rotation, une expérience calibrée à été effectuée, en effectuant un scan tomographique d’une colonne de différents types de grains dans des conditions semblables à celles des expériences triaxiales. Après ce premier scan, un

deuxième à été effectué après une rotation rigide de toute la colonne. Ceci implique un mouvement relatif nul entre les grains, une vecteur de déplacement différent pour tous les grains, mais *une rotation identique pour tous les grains*. La Figure 9.4 montre les deux images acquises, avec la rotation rigide de la colonne de 4.7° par rapport à l'axe verticale. Ces images on été traités avec les outils présentés plus haut, pour en venir à un suivi grain-par-grain.

La mesures de la rotations des grains fut abordé dans un premier moment en calculant la rotation d'une base orthonormale, définie avec les axes principales du tenseur d'inertie, définies dans la Section 5.3.4. Les rotations obtenues par ce biais ont un fort taux d'erreur: les lignes bleues dans les Figures 9.5 montrent une très grande distribution de rotations mesurées pour une rotation imposée de 4.7° . Ce résultat s'explique puisque à cette résolution, les eigenvecteurs du tenseur d'inertie de chaque grains ne sont pas suffisamment uniques, ce qui implique que la mesure n'est pas stable, et capable de donner des sauts de rotation mesurée importants.




Figure 9.5: Histogrammes montrant les distributions des rotations mesurés pour environ 9000 grains du sable d'Hostun et 6000 Ooids de Caicos, tous soumis à une rotation rigide de 4.7° . La largeur de corbeille est de 0.5°

La mesure de rotation a été abordée d'un autre angle: pour améliorer les qualités métrologiques de la mesure, de l'inspiration à été prise de la collaboration avec M. Bornert (publié dans Hall *et al.*, 2010): c'est à dire de la corrélation d'images *discrète*. Ceci veut dire qu'au lieu de faire de la corrélation d'images avec des sous-volumes de texture réguliers, les sous-volumes deviennent des

masque pour chaque grain identifié par l’algorithme de ligne de partage des eaux. À différence des l’approche développé dans Hall *et al.* (2010), la technique proposée ici fait un suivi discret avec ID-Track, et ensuite avec un grain déjà suivi, sa rotation est mesuré par une approche d’interpolation d’images. Ceci réduit énormément l’espace paramétrique de recherche, qui n’est pas pénalisé par le manque de degrés de libertés dans la première recherche, puisqu’elle est faite par suivi, et non pas par corrélation d’images. Cette technique à été publié Andò *et al.* (2012b), et l’algorithme est illustré synthétiquement dans la Figure 9.6. Des élément approfondis de l’algorithme peuvent de trouver dans la Section 5.3.4. Cette approche à la mesure de la rotation de grains donne une amélioration nette par rapport aux moments d’inertie: les lignes en rouge dans les Figures 9.5 montrent une erreur d’environ $\pm 1^\circ$ pour les grains d’Hostun (anguleux) et de $\pm 2^\circ$ pour les ooids de Caicos (qui sont beaucoup plus ronds).

9.3 Quelques résultats importants

Dans ce travail de thèse, les outils d’analyse discrète des images mentionnées plus haut (incluant un algorithme de comptage de contacts, *ID-Track* et la corrélation d’images discrète pour la mesure de rotation en 3D) ont été utilisés pour analyser les images 3D issues de la campagne expérimentale décrite plus haut. Ceci implique que pour chacune des trois sables étudiés, et pour deux pressions de confinement différentes, les cinématiques 3D de chaque grain (sauf un petite pourcentage de grains “perdus”) est mesuré entre tous les incrément acquis pendant l’essai.

Dans un premier moment ceci permet de vérifier que la cinématique des grains au tout début de chaque essai correspond à une déformation initialement homogène pour répondre à une attente intuitive du comportement des sols à des faibles déformations (bien avant une rupture ou une bifurcation dans le comportement). Ce genre de vérifications n’est pas normalement possible dans un essai triaxial classique, donc est loin de faire partie d’un standard d’opération, mais par contre ce genre de vérification peut aider à s’assurer que des échantillons petits tels que ceux étudiés se déforment de manière homogène. Un défaut géométrique de l’ordre de la taille d’un grain correspond à un très grand défaut de préparation quand rapporté à la taille d’un échantillon normal.

Prenant l’échantillon COEA01 comme exemple du genre d’analyse possible au sein de la campagne expérimentale, la Figure 9.7 montre dans un premier moment la réponse mécanique, macroscopique de cette éprouvette en compression triaxiale sous un confinement de 100 kPa. En haut on trouve R , le rapport de contrainte, et en bas la déformation volumique, les deux dessinées par rapport au taux de raccourcissement axial.

Cette réponse mécanique est ce qui est attendu d’un échantillon de sable dense: un pic dans la réponse en contrainte, avec un adoucissement vers un plateau de contrainte résiduelle après environ 9% de raccourcissement axial. D’un point de vue de déformation volumique, la réponse de l’échantillon COEA01 est aussi comme attendu pour un sable dense: avec une légère contraction tout au début du chargement, suivi par une dilatance importante de l’échantillon qui se stabilise quand la contrainte résiduelle est obtenue. La courbe supérieure est numérotée à 17 points distincts, qui correspondent aussi à une relaxation de la contrainte axiale – à ces points-ci, le chargement axial est interrompu pour faire un scan tomographique de l’échantillon.

Les Figures 9.8 et 9.9 montrent des coupes prises de l’essai COEA01, chacune orientée de la même manière par rapport à la tête de chargement supérieure de l’échantillon, qui reste immobile (le chargement axial se fait en comprimant l’échantillon du bas, vers le haut). Ces coupes sont orientées de telle manière à contenir à la fois l’axe verticale de l’échantillon, et le vecteur normale à la bande de cisaillement qui finit par se propager à travers l’échantillon à la fin de l’essai.

La Figure 9.8 montre des coupes dans un champs de porosité, mesuré de la manière mentionné plus haut. Cette mesure montre une distribution initialement homogène de la porosité de l’échantillon, qui ensuite se localise en une bande de porosité plus élevée traversant l’échantillon de manière diagonale. Cette mesure “mesoscopique” par rapport à la taille de l’échantillon est

Configuration Initiale

Configuration Déformée

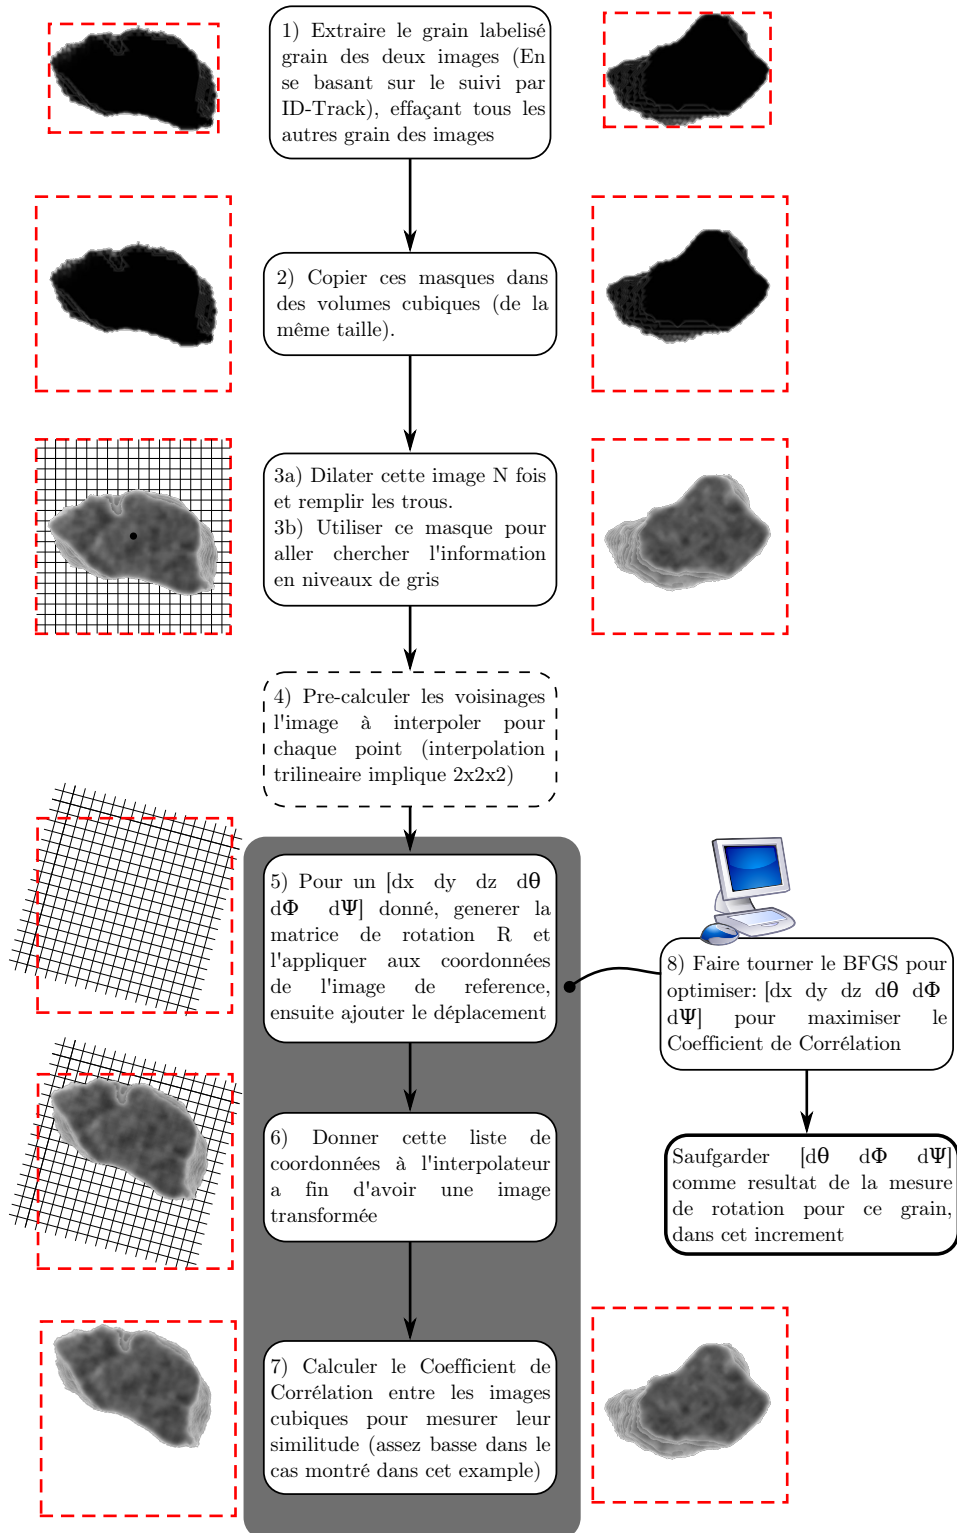


Figure 9.6: Diagramme d'opération montrant l'implémentation de la corrélation discrète 3D avec ID-Track, ce qui permet de mesurer la transformation rigide des grains. $N = 1$ dans cet travail

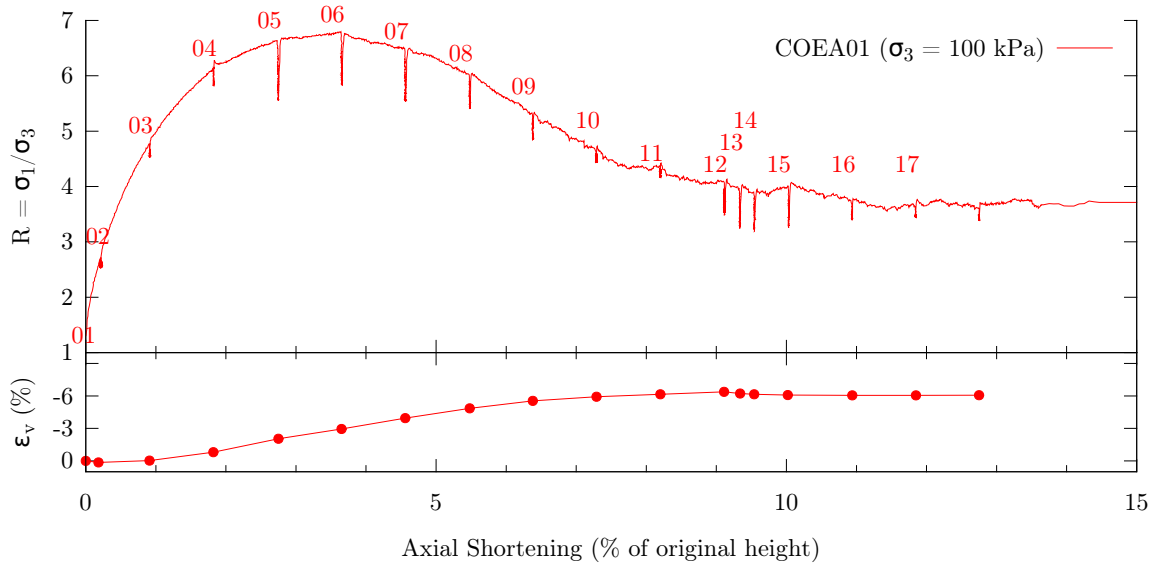


Figure 9.7: Réponse macroscopique de contrainte/déformation de l'échantillon COEA01




Figure 9.8: Des coupes dans une mesure d'un champs 3D porosité en différents états de l'essai COEA01

nommée de cette façon pour souligner que c'est une mesure qui reste *locale* par rapport aux mesures macroscopiques, mais reste à une échelle supérieure aux mesures à l'échelle du grain.

Les coupes dans la Figure 9.9 coupes sont un moyen pratique de montrer des quantités scalaires attachées aux grains – la première série d'images montre une quantité discrète attachée aux grains, calculée dans chaque image segmentée: le nombre de coordination de chaque grain (c'est-à-dire le nombre d'autres grain dont chaque grain est en contact – une information qui vient directement de l'algorithme de séparation d'objets, le ligne de partage d'eaux). L'évolution de cette quantité dans l'espace et dans le temps montre que au début de l'essai il y a une distribution qui semble (au moins selon la coupe) assez homogène de nombre de coordination,

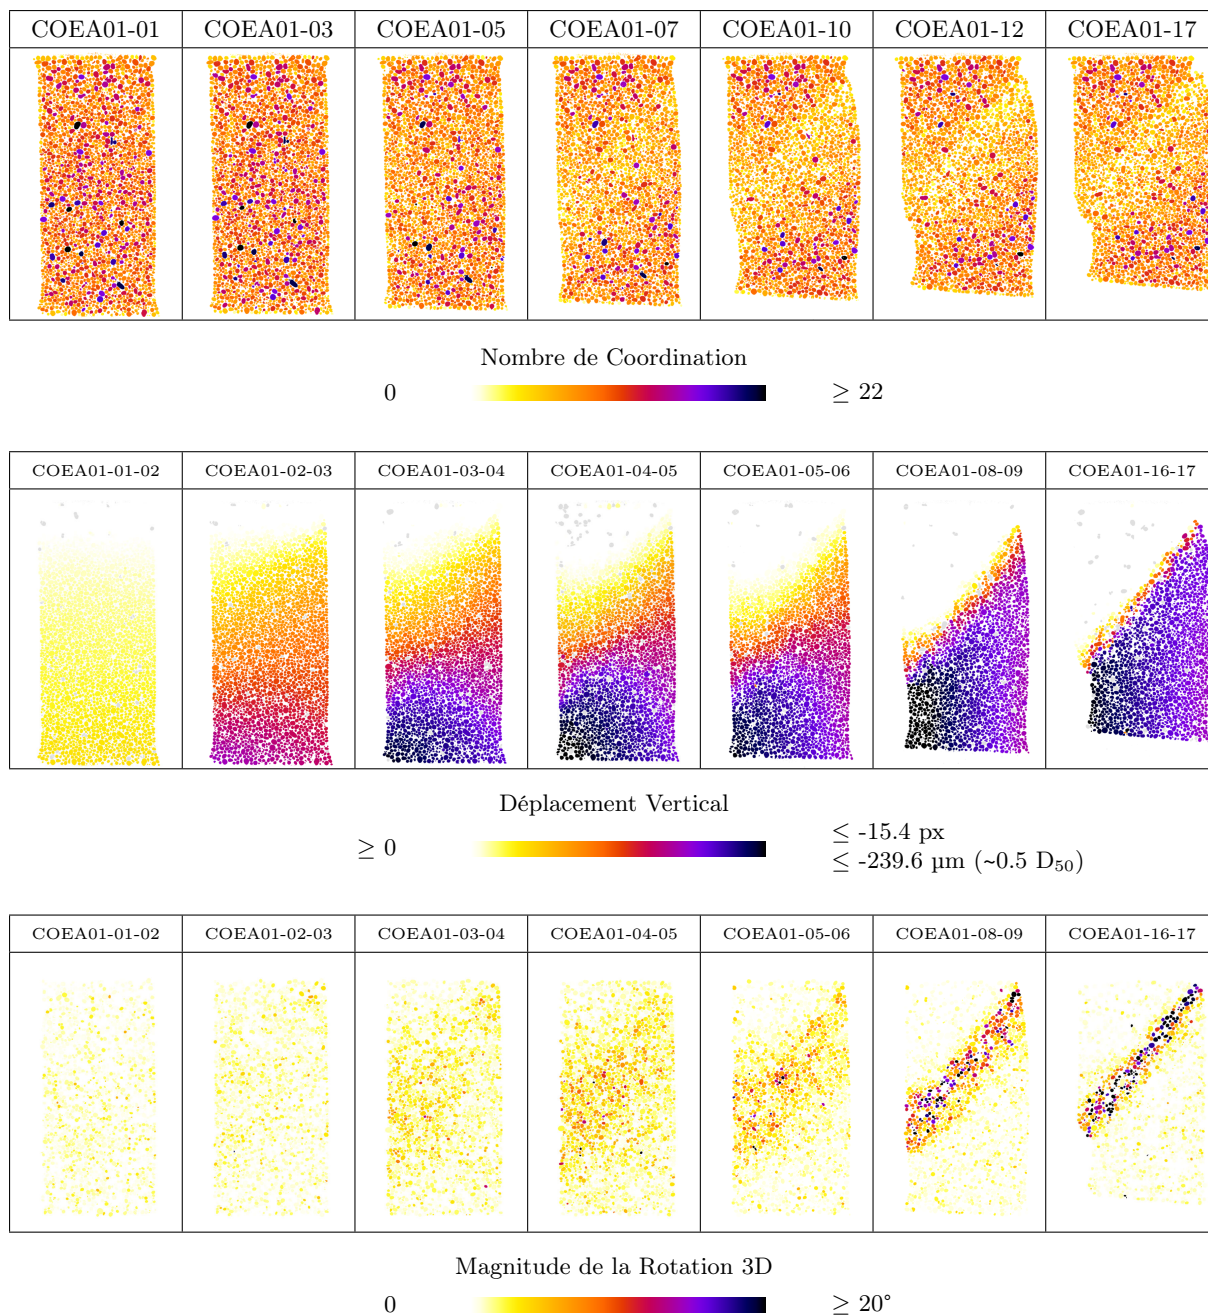


Figure 9.9: Certaines coupes verticales sélectionnées de l'essai COEA01. Chaque grain est coloré en fonction d'une quantité scalaire (dans ce cas le nombre de coordination de chaque grain, le déplacement vertical de chaque grain (un déplacement négatif = vers le haut), ou l'intensité de la rotation 3D de chaque grain), et une charte de couleur qui est donnée pour chaque scalaire. Pour les deuxièmes et troisièmes quantités (celles issues de ID-Track et de la corrélation discrète appliquée aux grain respectivement) les grains non-suivis sont montrés en gris clair.

mais étant loin d'une seule valeur pour tous les grains: en effet la grosse dispersion dans ces valeurs est un résultat de la distribution des tailles des grains qui forment l'échantillon: même si celle-ci est une distribution très serrée par rapport à des matériaux naturels, elle n'est pas non plus une distribution mono-disperse, ce qui a son tour veut dire que même un grain légèrement plus grand que ses voisins à la possibilité d'avoir quelques contacts de plus. L'évolution dans le temps de la distribution du nombre de coordination des grains dans cette coupe montre une perte de contacts dans une bande qui traverse l'échantillon du bas du côté gauche jusqu'en haut côté droit. Ceci est en correspondance avec une bande de cisaillement dilatante, ou la perte de

contacts des grains dans la bande correspond à une augmentation de volume, ou dilatance (ce qui est confirmé par l'augmentation de la porosité dans la bande vue ci-dessus).

Les lignes centrale et basse de la Figure 9.9 montrent des quantités attachées aux grains qui sont dérivées du suivi des grains d'un incrément au prochain. Les composantes verticales des vecteurs de déplacement calculé pour chaque grain montrent, au tout début de l'essai un champ de déplacement vertical axi-symétrique, ce qui correspond à un champ de déformation zz constant, chose qui est encourageante pour la qualité géométrique initial de l'échantillon COEA01, et sa réponse initiale. Après ce premier incrément, l'organisation spatiale du déplacement verticale des grains s'éloigne d'un champ axi-symétrique et subit une rotation progressive. Clairement visible aussi est la concentration du *gradient* de déplacement: dans le dernier incrément présenté, il est clair qu'il y a un déplacement nul dans la partie supérieure de l'échantillon, tandis que la partie inférieure à un champ déplacement qui correspond à une *rotation* rigide de la partie inférieure. La concentration principale de la déformation est dans une bande qui au milieu de l'échantillon, une bande de cisaillement. La palette de couleur ne facilite pas l'analyse de ces quantités, dû aux valeurs extrêmes aux bords du mécanisme à caractériser.

Les rotations des grains montrés sur la ligne basse de la Figure 9.9 montrent la partie scalaire de rotation quand une rotation tridimensionnelle est exprimée en tant que une axe de rotation et un scalaire de rotation autour de cet axe. Ces résultats montrent qu'au début de l'essai il a y pratiquement pas de rotations importantes. À la fin de l'essai la situation est aussi très simple: le haut et le bas de l'échantillon n'ont pas de rotations intenses, les rotations intenses sont toutes concentrées dans une bande d'une épaisseur de quelques grains. La transition entre ces deux états soi-disant "simples" est une transition complexe – l'apparition progressive d'une zone large (quand elle est comparée à la bande finale) ou les rotations sont concentrés se voit dès le deuxième incrément. Au pic de la contrainte macroscopique (incrément 05-06), la zone des rotation élevées est plus concentrée que dans le deuxième incrément, et on perçoit un nombre de grains en contact ou on pourrait avancer l'hypothèse qu'ils se transmettent une composante de leur rotation.

Pour mieux faire le lien entre ces mesures microscopiques, à l'échelle du grain, et les quantités plus usuelles de la mécanique des milieux continus, des quantités telles que la déformation peuvent être définies à micro-échelle, dans le style de ce qui est proposé dans Bagi (1996). Dans ce cas spécifique, les vecteurs de déplacement incrémentaux des grains sont insérés dans le programme d'Elements Discrets Yade Šmilauer *et al.* (2010), ce qui permet dans un premier moment de faire une tessellation 3D par tétraèdre de tous les centres des grains, et en insérant la cinématique, de calculer la déformation sur le volume du tétraèdre. Le deuxième invariant de ce champs (la déformation déviatoire) est montré pour l'essai COEA01 dans la Figure ??, à droite. Ce champs se compare bien aux rotations dans le même incrément montrés en bas, mais révèle aussi des concentrations dans le cisaillement dans la bande. Dans la même figure des champs similaires sont montrés pour un incrément dans l'état résiduel d'un essai sur les autres deux à la même pression de confinement (100 kPa). Étudiant les rotation en comparant à l'essai COEA01, c'est clair que le mécanisme de déformation dans la bande de cisaillement est d'un premier d'abord similaire. Une analyse plus en détail révèle que dans HNEA01 il y a une zone plus épaisse avec des rotations faibles (par rapport à la bande) mais non-nulles que dans COEA01 – ceci est expliqué par le fait que les grains beaucoup plus anguleux du sable d'Hostun ont beaucoup moins de liberté de tourner sans engendrer des rotations dans des grains voisins. Puisque la bande de cisaillement qui est bien mûre dans tous les cas de Figure 9.10, les gros déplacements imposés sur l'échantillon impliquent un fort taux de cisaillement, tout concentré dans la bande, qui pour déplacer les grains demande une rotation, qui donc se manifeste différemment d'un matériau à un autre. Le cas du sable d'Ottawa présente un cas intermédiaire entre les autres deux. La rotation moins libre des grains d'Hostun est certainement une raisons pour laquelle la contrainte résiduelle pour ce matériau est plus élevée que pour les autres deux, plus ronds.

L'analyse micro-mécanique de tous les essais est fait en détail dans le Chapitre 6.




Figure 9.10: Comparaison des mesures micro-mécaniques faites dans l'état résiduel des essais HNEA01 (sable d'Hostun), OUEA06 (sable d'Ottawa) et COEA01 (Caicos Ooids)

Le chapitre suivant (c'est-à-dire le Chapitre 7) montre des applications réelles du genre de donnée issue de ce travail dans le mode de la mécanique des milieux granulaires. Deux exemples sont choisis: le premier une simulation utilisation la méthode des éléments discrets, mais avec un arrangement de grains de sable réalistes. L'ensemble granulaire, en effet, est construit à partir d'un sous-volume dans un échantillon et la géométrie de chaque grain dans ce sous-volume est individuellement décrit au sein de la simulation. Le deuxième exemple de l'utilisation des résultats issues de cette travail est une analyse utilisant les réseaux complexes de l'ensemble de la cinématique de nombreux essais. Dans cette approche, des réseaux complexes sont créés en ignorant la position *spatiale* des grains, et en mettant en proximité seulement des grains avec une *cinématique similaire*. Ce réseau de points (*i.e.*, grains) dans cet espace est ensuite connecté par un critère de proximité (qui est défini comme la plus petite valeur de proximité qui permet d'avoir tous les points connectés dans un graphe). Une fois le réseau défini, des propriétés telles que la connectivité sont calculés, et peuvent être mises en jeu pour "détecter" avec le cinématique initiale l'endroit où se trouvera la bande à la fin de l'essai.

9.4 Conclusions

En conclusion dans ce travail de thèse, une campagne expérimentale d’essais de compression triaxiale sur trois différents types sables (d’angularité différente) à été réalisé sur des échantillons de taille réduite à fin de permettre la réalisation de ces essais *in-situ* dans un tomographe à rayons-x, ce qui permet de faire une image tridimensionnelle de l’échantillon à plusieurs reprises pendant l’essai. Un nombre d’outils ont été mis au point pour permettre l’analyse de ces images à l’échelle du grain, et aussi le suivi des grains entre les différents images acquises pendant un essai. Ces mesures micro-mécaniques ont été présentés pour les différents essais réalisés. La mise en relation de certaines mesures micro-mécaniques – par exemple la nature de la rotation des grains dans une bande de cisaillement mûre – avec des phénomènes macroscopiques – par exemple la contrainte résiduelle. Cette démarche micro-macro, qui est juste amorcée dans ce travail est très prometteuse vu les efforts de la communauté dans la direction de la micro-mécaniques, à la fois dans la modélisation explicite, telle que les éléments discrets, ou avec la prise en compte de paramètres micro-mécaniques (telles que les tenseurs de texture (“fabric tensors”) dans des modèles continus.

Bibliography

- Alshibli, K. and Alramahi, B. (2006), “Microscopic evaluation of strain distribution in granular materials during shear”, in: *Journal of geotechnical and geoenvironmental engineering*, Vol. 132, No. 1, pp. 80–91 (cit. on p. 16).
- Alshibli, K. A. and Reed, A. H., eds. (2010), *Advances in Computed Tomography for Geomaterials*, New Orleans, U.S.: Wiley (cit. on p. 16).
- Als-Nielsen, J. and McMorrow, D. (2011), *Elements of Modern X-ray Physics*, 2nd ed., Wiley and Sons (cit. on pp. 20, 22).
- Andò, E., Hall, S. A., Viggiani, G., Desrues, J., and Bésuelle, P. (2012a), “Grain-scale experimental investigation of localised deformation in sand: a discrete particle tracking approach”, in: *Acta Geotechnica*, Vol. 7, No. 1, pp. 1–13 (cit. on pp. 11, 81, 111, 117, 221).
- Andò, E., Hall, S. A., Viggiani, G., Desrues, J., and Bésuelle, P. (2012b), “Experimental micromechanics: grain-scale observation of sand deformation”, in: *Géotechnique Letters*, Vol. 2, No. 3, pp. 107–112 (cit. on pp. 11, 77, 97, 111, 117, 223).
- Andrade, J., Avila, C., Hall, S., Lenoir, N., and Viggiani, G. (2011), “Multiscale modeling and characterization of granular matter: From grain kinematics to continuum mechanics”, in: *Journal of the Mechanics and Physics of Solids*, Vol. 59, pp. 237–250 (cit. on p. 14).
- Andrade, J. E. and Avila, C. F. (2012), “Granular element method (GEM): linking inter-particle forces with macroscopic loading”, in: *Granular Matter*, Vol. 14, (1), pp. 51–61 (cit. on p. 216).
- Andrade, J. E., Lim, K.-W., Avila, C. F., and Vlahinić, I. (2012), “Granular element method for computational particle mechanics”, in: *Comput. Methods Appl. Mech. Engrg.* Vol. 241–244, pp. 262–274 (cit. on pp. 201, 202).
- Bagi, K. (1996), “Stress and strain in granular assemblies”, in: *Mechanics of Materials*, Vol. 20, No. 3, pp. 165–177 (cit. on pp. 146, 169, 210, 227).
- Barber, C. B., Dobkin, D. P., and Huhdanpaa, H. (1996), “The Quickhull algorithm for convex hulls”, in: *ACM TRANSACTIONS ON MATHEMATICAL SOFTWARE*, Vol. 22, No. 4, pp. 469–483 (cit. on p. 106).
- Bay, B. K. (2008), “Methods and applications of digital volume correlation”, in: *The Journal of Strain Analysis for Engineering Design*, Vol. 43, No. 8, pp. 745–760 (cit. on pp. 113, 132).
- Bay, B. K., Smith, T. S., Fyhrrie, D. P., and Saad, M. (1999), “Digital Volume Correlation: Three-dimensional Strain Mapping Using X-ray Tomography”, in: *Experimental Mechanics*, Vol. 39, No. 3, pp. 217–226 (cit. on pp. 113, 114).
- Bažant, Z., Belytschko, T., and Chang, T. (1984), “Continuum Theory for Strain-Softening”, in: *Journal of Engineering Mechanics*, Vol. 110, No. 12, pp. 1666–1692 (cit. on p. 13).
- Bernard, L., Fave, S., Noirfalise, E., and Saragaglia, A. (2011), *Visilog 7 Reference Guide*, Manual, Noesis S.A., Espace technologique, Route de l’orme, Batiment Mercury 2, 91190 Saint Aubin, France: Noesis S.A. (cit. on pp. 80, 81).
- Beucher, S. (1991), “The watershed transform applied to image segmentation”, in: *Proc. Pfefferkorn Conf. on Signal and Image Processing in Microscopy and Microanalysis*, Cambridge, UK, pp. 299–314 (cit. on p. 76).
- Beucher, S. and Lantuejoul, C. (1979), “Use of Watersheds in Contour Detection”, in: *International Workshop on image processing: Real-time image and Motion detection/estimation, Rennes, France* (cit. on p. 78).

- Bishop, A. w. and Henkel, D. (1962), *The measurement of soil properties in the triaxial test*, 2nd ed., Edward Arnold, London (cit. on pp. 42, 48, 52, 57).
- Bornert, M., Chaix, J. M., Doumalin, P., Dupré, J. C., Fournel, T., Jeulin, D., Maire, E., Moreaud, M., and Moulinec, H. (2004), “Mesure tridimensionnelle de champs cinématiques par imagerie volumique pour l’analyse des matériaux et des structures”, in: *Instrumn Measmt Metrology*, No. 4, pp. 43–88 (cit. on p. 116).
- Bousquet, H., Djeddid, M., Khazar, E., Flavigny, E., Hajd-Sadok, M., Lanier, J., Lefebvre, P., Meghachou, M., and Zitouni, Z. (1993), *Compilation des essais triaxiaux de révolution sur le sable d’Hostun RF*, Interal Laboratoire 3SR Circulation (cit. on p. 69).
- Butterfield, R., Harkness, R. M., and Andrews, K. Z. (1970), “A stereo-photogrammetric technique for measuring displacement fields”, in: *Géotechnique*, Vol. 20, No. 3, pp. 308–314 (cit. on p. 111).
- Calvetti, F., Combe, G., and Lanier, J. (1997), “Experimental micromechanical analysis of a 2D granular material: relation between structure evolution and loading path”, in: *Mechanics of Cohesive-Frictional Materials*, Vol. 2, pp. 121–163 (cit. on pp. 15, 216).
- Catalano, E., Chareyre, B., and Barthélémy, E. (2013), “Pore-scale modeling of fluid-particles interaction and emerging poromechanical effects”, in: *International Journal for Numerical and Analytical Methods in Geomechanics*, Submitted (cit. on p. 146).
- Chambon, R. (2011), “Instability and bifurcation for inelastic geomaterials”, in: *Advances in Bifurcation and Degradation in Geomaterials. Proceedings of the 9th International Workshop on Bifurcation and Degradation in Geomaterials*, ed. by S. Bonelli, C. Dascalu, and F. Nicot, 1st, Springer Series in Geomechanics and Geoengineering (cit. on p. 14).
- Chambon, R., Cailherie, D., and Matsuchima, T. (2001), “Plastic continuum with microstructure, local second gradient theories for geomaterials: localization studies”, in: *International Journal of Solids and Structures*, Vol. 38, pp. 8503–8527 (cit. on p. 14).
- Colliat-Dangus, J.-L. (1986), “Comportement des matériaux granulaires sous fortes contraintes; Influence de la nature minéralogique du matériau étudié”, PhD thesis, Institut National Polytechnique de Grenoble (cit. on p. 44).
- Combe, A.-L. (1998), “Comportement du sable d’Hostun S28 au triaxial axisymétrique. Comparaison avec le sable d’Hostun RF.”, MA thesis, Université Joseph Fourier (cit. on pp. 44, 58, 59, 69, 70).
- Coop, M. (1990), “The mechanics of uncemented carbonate sands”, in: *Géotechnique*, Vol. 40, No. 4 (cit. on p. 65).
- Cosserat, E. and Cosserat, F. (1909), *Théorie des Corps déformables*, A. Hermann et Fils, Paris (cit. on p. 14).
- Coupric, C., Grady, L. J., Najman, L., and Talbot, H. (2011), “Power watershed: A unifying graph-based optimization framework”, in: *IEEE Trans. Pattern Anal. Mach. Intell.* Vol. 33, No. 7, pp. 1384–1399 (cit. on p. 99).
- Delaunay, B. N. (1934), “Sur la sphère vide”, in: *Bulletin of Academy of Sciences of the USSR*, Vol. 6, pp. 793–800 (cit. on p. 106).
- Desrues, J. (1984), “La localisation de la déformation dans les matériaux granulaires”, PhD thesis, L’Université Scientifique et Médicale ET l’Institut National Polytechnique de Grenoble (cit. on p. 12).
- Desrues, J., Chambon, R., Mokni, M., and Mazerolle, F. (1996), “Void ratio evolution inside shear bands in triaxial sand specimens studied by computed tomography”, in: *Géotechnique*, Vol. 46, No. 3, pp. 529–546 (cit. on pp. 16, 69, 103, 104, 154, 155, 189).
- Desrues, J. and Viggiani, G. (2004), “Strain localization in sand: an overview of the experimental results obtained in Grenoble using stereophotogrammetry”, in: *International Journal for Numerical and Analytical Methods in Geomechanics*, Vol. 28, pp. 279–321 (cit. on pp. 15, 111).

- Desrues, J., Viggiani, G., and Bésuelle, P., eds. (2006), *Advances in Computed Tomography for Geomaterials*, Aussois, France: Wiley (cit. on p. 16).
- Durán, O., Kruyt, N., and Luding, S. (2010), “Analysis of three-dimensional micro-mechanical strain formulations for granular materials: Evaluation of accuracy”, in: *International Journal of Solids and Structures*, Vol. 47, pp. 251–260 (cit. on p. 146).
- Enomoto, T., Kawabeii, S., Tatsuoka, F., Benedetto, H. D., Hayashi, T., and Duttine, A. (2009), “Effects of Particle Characteristics on the Viscous Properties of Granular Materials in Shear”, in: Vol. 49, No. 1, pp. 25–49 (cit. on p. 66).
- Feldkamp, L. A., Davis, L. C., and Kress, J. W. (1984), “Practical cone-beam algorithm”, in: *J. Opt. Soc. Am. A*, Vol. 1, No. 6 (cit. on p. 30).
- Finno, R. J., Harris, W. W., Mooney, M. A., and Viggiani, G. (1997), “Shear bands in plane strain compression of loose sand”, in: *Géotechnique*, Vol. 47, No. 1, pp. 149–165 (cit. on p. 173).
- Flat Panel X-ray Imaging* (2004), available at:
<http://www.varian.com/media/xray/products/pdf/Flat%20Panel%20Xray%20Imaging%2011-11-04.pdf> (cit. on p. 27).
- Flavigny, E., Desrues, J., and Palayer, B. (1990), “NOTE TECHNIQUE le sable d’Hostun «RF»”, in: *Rev. Franç. Géotech.* No. 53 (cit. on pp. 43, 44).
- Frenkel, G., Blumenfeld, R., Grof, Z., and King, P. R. (2008), “Structural characterization and statistical properties of two-dimensional granular systems”, in: *Physical Review E*, Vol. 77, (cit. on p. 107).
- Frey, J., Chambon, R., and Dascalu, C. (2012), “A two-scale poromechanical model for cohesive rocks”, in: *Acta Geotechnica*, pp. 1–18 (cit. on p. 14).
- Fu, P. and Dafalias, Y. F. (2011), “Fabric evolution within shear bands of granular materials and its relation to critical state theory”, in: *International Journal for Numerical and Analytical Methods in Geomechanics*, Vol. 35, No. 18, pp. 1918–1948 (cit. on p. 14).
- Geers, M. G. D., De Borst, R., and Brekelmans, W. A. M. (1996), “Computing strain fields from discrete displacement fields in 2d-solids”, in: *Int. J. Solids Structures*, Vol. 33, No. 29, pp. 4293–4307 (cit. on p. 115).
- Germain, P. (1973), “La méthode des puissances virtuelles en mécanique des milieux continus, Première partie: Théorie du second gradient”, in: *J. Mécanique*, No. 12, pp. 235–274 (cit. on p. 14).
- Glasser, O. and Boveri, M. (1933), *Wilhelm Conrad Rontgen, and the early history of the Roentgen rays*, John Hale, London (cit. on p. 21).
- Glassner, A. S. (1993), *Graphics Gems*, Morgan Kaufmann (cit. on p. 86).
- Grady, L. (2006), “Random Walks for Image Segmentation”, in: *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 28, No. 11, pp. 1768–1783 (cit. on p. 99).
- Greve, J. W. (1967), *Industrial Metrology*, Manufacturing Engineering Series, American Society of Tool and Manufacturing Engineers (cit. on p. 86).
- Hall, S., Bornert, M., Desrues, J., Pannier, Y., Lenoir, N., Viggiani, G., and Bésuelle, P. (2010), “Discrete and continuum experimental study of localised deformation in Hostun sand under triaxial compression using X-ray μ CT and 3D digital image correlation”, in: *Géotechnique*, Vol. 60, No. 5, pp. 315–322 (cit. on pp. 16, 43, 52, 111, 116, 117, 125, 126, 128, 139, 146, 147, 222, 223).
- Hall, S., Lenoir, N., Pannier, Y., Desrues, J., Bornert, M., Viggiani, G., Bésuelle, P., Di Michiel, M., and Otani, J. (2008), “Characterisation of localised deformation in granular geomaterials using x-ray (micro) tomography and 3D-volumetric digital image correlation”, in: (cit. on p. 117).
- Hall, S. A., Wright, J., Pirling, T., Andò, E., Hughes, D. J., and Viggiani, G. (2011), “Can intergranular force transmission be identified in sand?, First results of spatially-resolved neutron and X-ray diffraction”, in: *Granular Matter*, Vol. 13, No. 3, pp. 251–254 (cit. on p. 217).

- Havelock, D. I. (1989), “Geometric Precision in Noise-Free Digital Images”, in: *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 11, No. 10, pp. 1065–1075 (cit. on pp. 89, 90).
- Havelock, D. I. (1991), “The Topology of Locales and Its Effects on Position Uncertainty”, in: *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 13, No. 4, pp. 380–386 (cit. on p. 90).
- Herman, G. T. and Davidi, R. (2008), “Image reconstruction from a small number of projections”, in: *Inverse Problems*, Vol. 24, No. 4 (cit. on p. 30).
- Hill, S. (1994), “Tri-linear Interpolation”, in: *Graphics Gems IV*, ed. by P. S. Heckbert, AP Professional (Academic Press), Boston, pp. 521–525 (cit. on p. 128).
- Hsieh, J. (2009), *Computed Tomography Principles, Design, Artifacts, and Recent Advances*, 2nd ed., Wiley (cit. on pp. 21, 22, 29, 30).
- Hubbell, J. H. and Seltzer, S. M. (1996), *Tables of X-Ray Mass Attenuation Coefficients and Mass Energy-Absorption Coefficients from 1 keV to 20 MeV for Elements Z = 1 to 92 and 48 Additional Substances of Dosimetric Interest*, available at: <http://physics.nist.gov/PhysRefData/XrayMassCoef/cover.html>, NIST 5632 Report, (NIST Standard Reference Database 126) (cit. on p. 23).
- Jaquet, C., Andò, E., Viggiani, G., and Talbot, H. (2013), “Estimation of separating planes between touching 3D objects using power watershed”, in: Vol. 7883, pp. 452–463 (cit. on pp. 100, 210).
- Jolliffe, I. T. (2002), *Principal Component Analysis*, Second, Springer, ISBN: 0387954422 (cit. on p. 99).
- Kaddhour, G., Ando, E., Salager, S., Bésuelle, P., Viggiani, C., Hall, S., and Desrues, J. (2013), “Application of X-ray Tomography to the Characterisation of Grain-Scale Mechanisms in Sand”, in: *Multiphysical Testing of Soils and Shales*, ed. by L. Laloui and A. Ferrari, Springer Series in Geomechanics and Geoengineering, Springer Berlin Heidelberg, pp. 195–200, ISBN: 978-3-642-32491-8, DOI: 10.1007/978-3-642-32492-5_23, URL: http://dx.doi.org/10.1007/978-3-642-32492-5_23 (cit. on pp. 104, 216).
- Khaddour, G. (2012), “Microscopic Mechanisms Related to Water Retention in Sand”, MA thesis, Université de Grenoble (cit. on p. 216).
- Kim, H.-K. and Santamarina, J. (2008), “Sand–rubber mixtures (large rubber chips)”, in: *Canadian Geotechnical Journal*, Vol. 45, No. 10, pp. 1457–1466 (cit. on pp. 44, 46).
- Koks, D. (2006), “A Roundabout Route to Geometric Algebra”, in: *Explorations in Mathematical Physics*, Springer Science+Business Media, LLC., pp. 147– (cit. on pp. 124, 141).
- Kouznetsova, V., Brekelmans, W. A. M., and Baaijens, F. P. T. (2001), “An approach to micro-macro modeling of heterogeneous materials”, in: *Computational Mechanics*, Vol. 27, pp. 37–48 (cit. on p. 14).
- Landau, U. M. (1987), “Estimation of a Circular Arc Center and Its Radius”, in: *Computer Vision, Graphics, and Image Processing*, Vol. 38, pp. 317–326 (cit. on p. 89).
- Lee, J., Dodds, J., and Santamarina, J. (2007), “Behavior of rigid-soft particle mixtures”, in: *Journal of Materials in Civil Engineering*, Vol. 19, No. 2, pp. 179–184 (cit. on p. 62).
- Lenoir, N., Bornert, M., Desrues, J., Bésuelle, P., and Viggiani, G. (2007), “Volumetric Digital Image Correlation Applied to X-ray Microtomography Images from Triaxial Compression Tests on Argillaceous Rock”, in: *Strain*, Vol. 43, pp. 193–205 (cit. on pp. 52, 114).
- Lenoir, N. (2006), “Comportement mécanique et rupture dans les roches argileuses étudiés par micro tomographie à rayons X”, available at: <http://tel.ccsd.cnrs.fr/tel-00011996>, PhD thesis, University of Grenoble, France (cit. on p. 51).
- Lesniewska, D. and Wood, D. M. (2011), “Photoelastic and photographic study of a granular material”, in: *Géotechnique*, Vol. 61, No. 7, pp. 605–611 (cit. on p. 15).

- Lorensen, W. and Cline, H. (1987), “Marching cubes: A high resolution 3D surface construction algorithm”, in: *Computer Graphics*, Vol. 21, pp. 163–169 (cit. on pp. 93, 95).
- Majmudar, T. S. and Behringer, R. P. (2005), “Contact force measurements and stress-induced anisotropy in granular materials”, in: *Nature*, Vol. 435, pp. 1079–1082 (cit. on p. 15).
- Marak, L., Couprie, M., and Talbot, H. (2011), “The Pink Image Processing Library”, in: *Euroscipy - European conference on scientific python 2011*, <http://www.pinkhq.com>, <http://www.euroscipy.org/talk/4190> (cit. on p. 97).
- Matsushima, T., Saomoto, H., Uesugi, K., Tsuchiyama, A., and Nakano, T. (2004), “Detection of 3-D irregular grain shape of Toyoura sand at SPring-8, X-ray CT for Geomaterials”, in: *Proc. International workshop on X-ray CT for geomaterials*, ed. by Otani and Obara, pp. 121–126 (cit. on p. 77).
- Matsushima, T., Uesugi, K., Nakano, T., and Tsuchiyama, A. (2006), “Visualization of grain motion inside a triaxial specimen by micro X-ray CT at SPring-8”, in: *Advances in X-ray Tomography for Geomaterials*, pp. 35–52 (cit. on p. 16).
- Matsushima, T., Chambon, R., and Caillerie, D. (2002), “Large strain finite element analysis of a local second gradient model: application to localization”, in: *International Journal for Numerical Methods in Engineering*, Vol. 54, pp. 499–521 (cit. on p. 14).
- Meyer, F. (1994), “Topographic distance and watershed lines”, in: *Signal processing*, Vol. 38, No. 1, pp. 113–125 (cit. on p. 79).
- Mortensen Montoya, B. (2012), “Bio-Mediated Soil Improvement and the Effect of Cementation on the Behavior, Improvement, and Performance of Sand”, PhD thesis, University of California Davis (cit. on p. 44).
- Nakamura, A. and Aizawa, K. (1984), “Digital Circles”, in: *Computer Vision, Graphics, and Image Processing*, Vol. 26, pp. 242–255 (cit. on p. 89).
- Nitka, M., Combe, G., Dascalu, C., and Desrues, J. (2011), “Two-scale modeling of granular materials: a DEM-FEM approach”, in: *Granular Matter*, Vol. 13, No. 3, pp. 277–281 (cit. on pp. 14, 217).
- Novotni, M. and Klein, R. (2001), “A geometric approach to 3D object comparison”, in: *Shape Modeling and Applications, SMI 2001 International Conference on*. Pp. 167–175 (cit. on p. 125).
- Oda, M., Takemura, T., and Takahashi, M. (2004), “Microstructure in shear band observed by microfocus X-ray computed tomography”, in: *Géotechnique*, Vol. 54, No. 8, pp. 539–542 (cit. on p. 13).
- O’Gorman, L. (1991), “The Topology of Locales and Its Effects on Position Uncertainty”, in: *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 18, No. 7, pp. 746–751 (cit. on p. 89).
- Osher, S. and Fedkiw, R. (2003), “Level Set Methods and Dynamic Implicit Surfaces”, in: *Applied Mathematical Sciences*, Vol. 153, (cit. on pp. 81, 200).
- O’Sullivan, C. (2012), “Particle-Based Discrete Element Modeling: Geomechanics Perspective”, in: *International Journal of Geomechanics*, Vol. 11, No. Special Issue: Material and Computer Modeling, pp. 449–464 (cit. on p. 15).
- Otani, J. and Obara, Y., eds. (2003), *X-ray CT for Geomaterials; Soils, Concrete, Rocks*, Kumamoto, Japan: A.A. Balkema (cit. on p. 16).
- Otsu, N. (1975), “A threshold selection method from gray-level histograms”, in: *Automatica*, Vol. 11, No. 285-296, pp. 23–27 (cit. on p. 74).
- Pan, B., Qian, K., Xie, H., and Asundi, A. (2009), “Two-dimensional digital image correlation for in-plane displacement and strain measurement: a review”, in: *Measurement Science and Technology*, Vol. 20, No. 6 (cit. on p. 112).

- Pannier, Y., Lenoir, N., and Bornert, M. (2010), “Discrete volumetric digital image correlation for the investigation of granular type media at microscale: accuracy assessment”, in: *EPJ Web of Conferences*, Vol. 6, p. 35003 (cit. on pp. 133, 135).
- PaxScan® 2520XI/V* (2010), available at: www.varian.com/media/xray/products/pdf/ps2520v.pdf (cit. on p. 27).
- Piotrowska, E. (2013 - expected), “Rôle du squelette granulaire dans le comportement du béton sous très fortes contraintes : analyse expérimentale et numérique.”, PhD thesis, Université de Grenoble (cit. on p. 138).
- Radon, J. (1917), “Über die Bestimmung von Funktionen durch ihre Integralwerte Langs Gewisser Mannigfaltigkeiten (English translation: On the determination of functions from their integrals along certain manifolds)”, in: *Ber. Saechsische Akad. Wiss.* Vol. 29, p. 262 (cit. on p. 29).
- Riedel, I. (2011), “Characterisation and Micro-Scale Analysis of Hostun Sand Water Retention Behaviour”, MA thesis, Université de Grenoble (cit. on p. 216).
- Roscoe, K. H. and Burland, J. B. (1968), “On the Generalized Stress-Strain Behaviour of Wet Clay”, in: *Engineering plasticity*, pp. 535–609 (cit. on p. 12).
- Saadatfar, M., Sheppard, A. P., Senden, T. J., and Kabla, A. J. (2012), “Mapping forces in a 3D elastic assembly of grains”, in: *Journal of the Mechanics and Physics of Solids*, Vol. 60, No. 1, pp. 55–66, DOI: 10.1016/j.jmps.2011.10.001 (cit. on p. 217).
- Satake, M. (1989), “Mechanics of Granular Materials, Report of ISSMFE Technical Committee, Rio de Janeiro”, in: Japanese Society of Soil Mechanics and Foundation Engineering (cit. on p. 146).
- Seibert, J. A., Boone, J. M., and Lindfors, K. K. (1998), “Flat-field correction technique for digital detectors”, in: pp. 348–354, DOI: 10.1117/12.317034, URL: [+%20http://dx.doi.org/10.1117/12.317034](http://dx.doi.org/10.1117/12.317034) (cit. on p. 27).
- Sibille, L. and Froiio, F. (2007), “A numerical photogrammetry technique for measuring microscale kinematics and fabric in schneebeli materials”, in: *Granular Matter*, Vol. 9, pp. 183–193 (cit. on pp. 15, 160, 161).
- Sibleco France (2011), *Fiche Technique Type HN31*, available at: http://www.sibelco.fr/item_img/medias/images/ft12_hn31.pdf (cit. on pp. 43, 46).
- Šmilauer, V., Catalano, E., Chareyre, B., Dorofenko, S., Duriez, J., Gladky, A., Kozicki, J., Modenese, C., Scholtès, L., Sibille, L., Stránský, J., and Thoeni, K. (2010), “Yade Reference Documentation”, in: *Yade Documentation*, ed. by V. Šmilauer, 1st, available at: <http://yade-dem.org/doc/>, The Yade Project (cit. on pp. 146, 169, 227).
- Smith, T., Bay, B., and Rashid, M. (2002), “Digital Volume Correlation Including Rotational Degrees of Freedom during Minimization”, in: *Experimental Mechanics*, Vol. 42, No. 3, pp. 272–278 (cit. on pp. 113, 133).
- Tagliaferri, F., Waller, J., Andò, E., Hall, S. A., Viggiani, G., Bésuelle, P., and DeJong, J. T. (2011), “Observing strain localisation processes in bio-cemented sand using x-ray imaging”, in: *Granular Matter*, Vol. 13, No. 3, pp. 247–250 (cit. on pp. 104, 215).
- Terzaghi, K. von (1920), “Old Earth Pressure Theories and New Test Results”, in: *Engineering News-Record*, Vol. 85, No. 14, pp. 632–637 (cit. on p. 13).
- Tordesillas, A., Walker, D. M., Andò, E., and Viggiani, G. (2013), “Revisiting localised deformation in sand with complex systems”, in: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, Vol. 469, No. 2152 (cit. on pp. 202, 204, 206–208).
- Truesdell, C. and Noll, W. (1965), *The Non-Linear Field Theories of Mechanics*, Springer (cit. on pp. 13, 14).
- Vardoulakis, I. and Sulem, J. (1995), *Bifurcation Analysis in Geomechanics*, Blackie Academic and Professional (cit. on p. 13).
- Viggiani, G., Küntz, M., and Desrues, J. (2001), “An experimental investigation of the relationships between grain size distribution and shear banding in sand”, in: *Continuous and*

- Discontinuous Modelling of Cohesive-Frictional Materials*, ed. by P. Vermeer, H. Herrmann, S. Luding, W. Ehlers, S. Diebels, and E. Ramm, vol. 568, Lecture Notes in Physics, Springer Berlin Heidelberg, pp. 111–127, ISBN: 978-3-540-41525-1 (cit. on p. 214).
- Vlahinić, I., Andò, E., Viggiani, G., and Andrade, J. E. (2012), “Towards a more accurate characterization of granular media: extracting quantitative descriptors from grain-scale images”, in: *Granular Matter*, Submitted (cit. on pp. 81, 90, 200, 201).
- Walker, D. M. and Tordesillas, A. (2012), “Taxonomy of granular rheology from grain property networks”, in: *Phys. Rev. E*, Vol. 85, (cit. on p. 202).
- Wicher, K. (2012), “Technique to calculate of the evolution of void ratio from tomographic images of different granular materials subjected to isotropic compression”, MA thesis, Université Joseph Fourier (cit. on p. 107).
- Wiebicke, M. (2012), “Accessing Forces between Elastic Grains using 2D Digital Image Correlation”, MA thesis, Université de Grenoble (cit. on p. 216).
- Yoo, T., Ackerman, M., Lorensen, W., Schroeder, W., Chalana, V., Aylward, S., Metaxes, D., and Whitaker, R. (2002), “Engineering and Algorithm Design for an Image Processing API: A Technical Report on ITK - The Insight Toolkit”, in: *Proc. of Medicine Meets Virtual Reality*, ed. by J. Westwood, pp. 586–592 (cit. on p. 116).