

Chapter 5

Multiscale Characterization and Modeling of Granular Materials through a Computational Mechanics Avatar: A Case Study with Experiment

5.1 Introduction

In recent times, much of solid mechanics research has focused around the ‘microstructures’ theme, which is predicated on the importance of lower-scale geometry, defects, and interaction to inform macroscopic behavior [114]. Crystalline solids have received significant attention, and progress in theory, experiment, and computation has been made to the extent that much of the research results can now be routinely applied in the design of such materials for real engineering applications. Granular materials also possess microstructures, which have far-reaching implications on macroscopic properties such as strength, permeability, and energy dissipation. Unlike crystalline solids, however, microstructures in granular materials lack order and cannot be elegantly quantified as in crystalline solids. Moreover, these microstructures operate at the grain scale, which neither atomistic nor continuum models can help in providing quantitatively meaningful results for real applications. To understand and predict the behavior of granular materials at the continuum scale, one must recognize that their mechanical behavior is encoded at the grain scale. We often point to this fact, despite the common knowledge that more than three decades of research has not led to significant advances in terms of performing grain-scale analysis and characterization

of real granular materials.

There are broadly two camps in the study of granular materials at the grain scale: that which seeks to characterize granular assemblies, often through imaging techniques such as X-ray computed tomography (XRCT), and that which seeks to simulate granular behavior, typically through the discrete element method (DEM). Characterization of granular materials has come a long way since the days of so-called ‘destructive’ methods, which involved dismantling experimental specimens to measure grain-scale quantities [115]. XRCT and other methods of digital imaging of experiments are now the norm in experimental grain-scale mechanics [17; 116], able to measure quantities such as the void ratio in a shear band [117] and track the motion of grains [118]. Increases in image fidelity and resolution have led to a number of characterization approaches such as level sets [47], spherical harmonics [119; 120] and Fourier analyses [121; 122] that can process image data to quantify grain kinematics and morphological measures (e.g., sphericity and roundness) [1]. Although characterization techniques have seen tremendous progress, they lack the ability to probe one crucial aspect of granular materials from which strength-related quantities are derived: interparticle contact forces. As such, a discrete model such as DEM remains a necessary component for the inference of contact forces. In the several decades following the inception of DEM [24], there were tremendous efforts in the development of shape representation capabilities and associated contact algorithms (e.g., [35–37; 40]). Currently, however, it appears that progress has hit a plateau with shape representations, largely belonging to either the polyhedra or clustering approach, still too crude for real grain-scale calculations. Incidentally, grain morphological measures from characterization rarely make their way into DEM because of incompatibilities between the geometry bases used in characterization and discrete models, and significant degradation of fidelity during the image-to-model process.

Despite great research efforts in each of the two microscopic camps, there is a surprising lack of interaction, and hence integration and validation, between them. In theory, both grain-scale characterization and discrete simulations should work in an integrated manner to provide a consistent set of microscopic information. In practice, however, this consistency is virtually nonexistent and surprisingly, not discussed in the literature. The reason for this void is that much of the problems making integration and validation difficult stem from the complexity of real grain geometries. For example, we described in an earlier work [22] two possible routes to using grain-scale information at the continuum level by means of a

hierarchical multiscale scheme: one using XRCT data and the other using discrete analysis. The consistency between the two routes, however, could not be established at the time since there was no DEM technology to account for the level of complexity of grain geometries found in experiment.

At the other end of the spectrum, multiscale methods are emerging to enable the use of grain-scale information at the continuum level. At this point, techniques linking the grain and continuum scales are based on homogenization theory [23]. It is important to point out that multiscale modeling is not just about developing algorithms but also about developing better physical models [123]. In the context of real granular materials, the effects of grain morphology on strength of granular materials have already been well established [1]. Although there have been significant efforts in constructing multiscale procedures (e.g., [124–131]), the predictive capabilities of these procedures on real granular materials have not been assessed due to the absence of any discrete model that is morphologically representative of real granular materials. Currently, the fidelity of the discrete model appears to be the missing ingredient to achieving a breakthrough in the predictive power of multiscale methods.

Here, we tackle the above challenges head-on by developing an overarching computational mechanics avatar that has enabled us to make the first steps in bridging the gap between characterization and discrete analysis, and potentially improving the predictive capabilities of multiscale methods. We integrate two major components of the proposed avatar, namely the level-set characterization technique and our NURBS-based DEM, and significantly enhance their capabilities to handle real grain morphologies. We then apply the avatar to characterize and model the grain-scale response inside the shear band of a real triaxial specimen. The avatar has enabled, for the first time, the transition from faithful representation of grain morphologies in X-ray tomograms of granular media to a morphologically accurate discrete computational model. Grain-scale information is extracted and upscaled into a continuum finite element model through a hierarchical multiscale scheme, and the onset and evolution of a persistent shear band is modeled, showing excellent quantitative agreement with experiment in terms of both grain-scale and continuum responses in the post-bifurcation regime. More importantly, consistency in results across characterization, discrete analysis, and continuum response from multiscale calculations is found, achieving the first and long sought-after quantitative breakthrough in grain-scale analysis

of real granular materials.

This chapter is organized as follows. First, we describe the characterization and computational components of the avatar, and the associated algorithmic improvements that enable faithful representation of grain morphologies directly from X-ray tomograms. Next, we present a case study where we describe the experiment of interest, followed by the characterization and modeling steps taken to arrive at the relevant grain-scale information and predicted continuum response. Finally, we close by providing an assessment of the strengths and validity of the components of the avatar, as well as limitations that are to be addressed in the future.

5.2 The Avatar Framework

A computational avatar aims to be the virtual twin of a granular assembly (e.g., sand), a digital mirror of grains as they are found in-situ. The avatar should also possess realistic mechanical properties (e.g., elasticity) and in a mechanistic context, behave similarly to the real grain assembly, so that grain morphologies, kinematics, arrangement, and elastostatics are all comparable to those of the physical assembly. The avatar opens the door for the use of the experimentally-derived grain-scale quantities in either forward simulation techniques (e.g., DEM) or inverse elastostatics analyses [2; 3]. These analyses would yield the quantitative measurements of contact forces throughout a granular assembly, allowing access to force chains or fabric evolution in real time in real granular assemblies.

In this section, we summarize the two major components of the proposed computational mechanics avatar that has made the image-to-model transition possible, namely the characterization toolbox used to capture grain morphologies and our NURBS-based DEM for discrete analysis. In particular, we focus on the relevant algorithmic improvements. Earlier, these components have been briefly discussed in [4] and developed independently [47; 49]. Here, we integrate these two components for the first time to enable the application of the avatar to a real problem.

5.2.1 Characterization toolbox

Successful inference of kinematics and contacts provides several important applications concerning the evolution of strength in granular systems, namely:

1. Determination of grain kinematics and grain fabric, including contact evolution during loading.
2. Inference of contacts locations. Together with strains or stresses from X-ray diffraction [19], contact locations provide necessary input for a technique for quantitative inference of contact forces in opaque granular systems [2; 3].
3. Inference of grain morphologies that are representative of true grains (to within imaging resolution). These serve as geometry input to discrete methods that can account for arbitrary grain shapes (see Section 5.2.2).

The use of XRCT in an experiment (e.g., triaxial test) provides, in addition to macroscopic stress and strain data, a sequence of three-dimensional (3D) of voxelated images (see Figure 5.1), each containing tens of thousands of grains, collected over the course of the specified loading regime. The challenge in applying XRCT is in translating the 3D images into quantities that can be used for mechanical analysis, i.e., to translate image voxels into grain fabric and morphology. The key difficulty underlying this translation process lies in the irregularity of individual grains and that they are in contact with each other.




Figure 5.1: Grain-scale imaging using XRCT: (a) slice of triaxial specimen and (b) 3D reconstruction by stacking slices.

The watershed technique [132] has been a trusted workhorse for segmenting grains from the voids, and from each other. Watershed, however, has the drawback of generally oper-

ating on and outputting binary images. This is problematic for two reasons. The first is that binary images introduce artificial roughness to grain surfaces, complicating a direct tomography-to-simulation paradigm [4]. The second, and more critical, drawback is the removal of details about the location and orientation of interparticle contact, which impedes our understanding of the physical sources of mechanical strength.

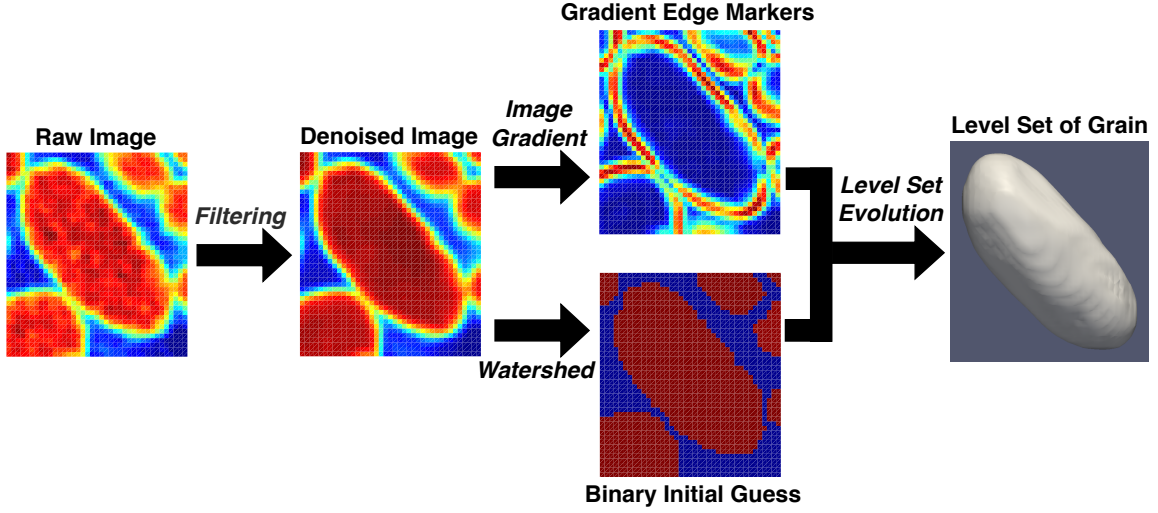


Figure 5.2: Summary of characterization steps: (left) slice of 3D XRCT volume of voxels, (middle two) application of proposed characterization technique, and (right) generated grain.

We overcome the above difficulties by applying the characterization methodology proposed in [47] to operate on the following key areas to delineate grain surfaces and contact locations. Briefly, the steps in the proposed methodology are:

1. Use of full-fidelity 3D XRCT image instead of binary images. The image is filtered to furnish precise edge markers via the first and second gradients of X-ray attenuation.
2. Noise removal. The use of image gradients necessitates the removal of noise. Non-local de-noising is utilized and guided by input parameters across different materials and X-ray attenuation spectra.
3. The search for grain edges via level sets. Speed up in convergence is obtained using the edge markers and current segmentation techniques (e.g., watershed) as initial conditions. The level sets are allowed to evolve until the boundary of grain is identified and, consequently, mathematically characterized.

The above steps produce grain boundaries that are smooth and representative of true grain shapes to sub-voxel accuracy, and without ‘melt’ near contact, as shown in Figure 5.2. In essence, the proposed methodology converts digital computed tomographic images into a collection of grain avatars, i.e., digitized versions of natural grains, which are transferred into a discrete computational model, as described next.

5.2.2 NURBS-based discrete element method

We employ a discrete element method that uses Non-Uniform Rational Basis-Splines (NURBS) as the mathematical basis for representing grain geometries [49]. The initial developments of this approach focused on strictly convex grains [49–51]. The generation of strictly convex NURBS shapes, however, is very difficult and restrictive from a modeling perspective. This is even more so when dealing with image data of real grain shapes and obtaining strictly convex shapes is not possible in most cases. This difficulty is not due to some limitation of NURBS but to the state of the contact algorithm for NURBS, which were undeveloped at the time.

A contact algorithm capable of dealing with general non-convex NURBS surfaces would eliminate the above modeling difficulty. We achieve this by generalizing the node-to-surface approach typically used in the contact treatment of finite element models [92] to a knot-to-surface approach. Following a master-slave approach, a fixed contact point on the slave particle Ω^j is denoted by \mathbf{x} while the contact point on the master particle Ω^i is defined to be the closest point projection of \mathbf{x} onto the boundary of the master particle:

$$\bar{\mathbf{y}} \equiv \bar{\mathbf{y}}(\mathbf{x}) = \min_{\mathbf{y} \in \Gamma^i} \|\mathbf{x} - \mathbf{y}\| \quad (5.1)$$

where Γ^i is the boundary of the master particle. Here, \mathbf{y} is a function of the two variables parametrizing the NURBS surface. Knowing the normal \mathbf{n} at $\bar{\mathbf{y}}$, the gap is then calculated as:

$$g(\mathbf{x}) = (\mathbf{x} - \bar{\mathbf{y}}(\mathbf{x}))^T \mathbf{n} \quad (5.2)$$

from which the effective normal contact force on grain Ω^i at $\bar{\mathbf{y}}$, assuming a linear elastic

stiffness model, is calculated:

$$\mathbf{f}_n^i = \begin{cases} k_N g \mathbf{n}, & \text{if } g < 0 \\ \mathbf{0}, & \text{otherwise} \end{cases} \quad (5.3)$$

where k_N is the normal elastic stiffness and $g \mathbf{n}$ measures the penetration or overlap in vector form.

The above approach is described in [52] for the two-dimensional (2D) case and here, we generalize this to 3D as follows. As a pre-processing step, we first perform a least-squares NURBS fitting procedure on the level set surfaces generated from the characterization step. Then, we apply a NURBS recursive subdivision procedure [50] to generate a set of fixed knots or nodes on each grain. The NURBS surface is subdivided until the following termination criteria (see Figure 3.3) are met:

1. Real space arc distance corresponding to each parametric direction of less than ϵ_d
2. Acute angle between normal vectors at adjacent knot positions of less than ϵ_θ

and a knot or node is positioned at the center of the subdivided surface at the terminated level (see Figure 3.3).

During simulation, we take advantage of the parametric nature of NURBS and solve the 3D closest point projection problem (5.1) using a two-stage optimization strategy: the proximity of the location of the closest projected point is established, followed by a local constrained optimization step to finalize the closest point location. The first optimization stage is handled using the DIviding RECTangles (DIRECT) global optimization algorithm [94], which effectively deals with non-convex objective functions and hence, non-convex grain shapes, while the second stage is handled using standard derivative-free optimization techniques [98]. This two-stage strategy eliminates all the major problems associated with intersection-based approaches to computing the interparticle gap or overlap (high computational cost, data proliferation, and convexity restriction) in the early version of our NURBS-based DEM [50]. Moreover, this procedure can be formulated and implemented with equal ease as in the 2D case.

5.3 Case Study

The purpose of the case study is to apply and assess the predictive capabilities of the avatar in a real experiment. This is achieved as follows. First, the localized incremental displacement fields and global stress-strain curve are extracted from experiment and used in a hierarchical multiscale computation to reproduce the continuum response. In this process, we obtain two quantities of interest in the shear band, namely the average dilatancy evolution and residual shear strength. Then, unit cells in the shear band are modeled using NURBS, and the same two quantities above are extracted and compared with those inferred from experiment. Below, we describe the experiment setup followed by the characterization and modeling steps leading to the final assessment.

5.3.1 Experiment setup

A cylindrical specimen (11 mm diameter by 24 mm height) of Caicos ooids (a natural granular material with well-rounded grains) was sheared at a constant axial strain rate of 0.1%/min under a constant cell pressure of 100 kPa in a specially designed triaxial cell in Laboratoire 3SR [133]. Tomographic images were captured at 18 stations during the test with a voxel size of $15.56^3 \mu\text{m}^3$, to be compared to a mean grain diameter $420 \mu\text{m}$. The global deviator stress as a function of the axial strain is shown in Figure 5.3. The behavior is typical of a dense sand, with a peak in the deviator stress followed by softening and subsequent residual or critical state and dilatant behavior throughout. The progressive formation of a shear band was observed during the test, with a final shear band inclination of approximately 52° with respect to the horizontal, and a thickness of approximately 8 grains.

5.3.2 Dilatancy inference

The dilatancy evolution in the shear band is inferred from grain-scale kinematics. Digital image correlation (DIC) is used in concert with tomographic images to compute incremental displacement vectors in elements containing 3^3 grains on average. The result is a 3D incremental displacement field for the entire triaxial specimen at all tomographic measurement stations. For example, Figure 5.4 shows the incremental displacement field between tomographic stages 6 and 7 (post-peak in the load-displacement curve). Some large incre-

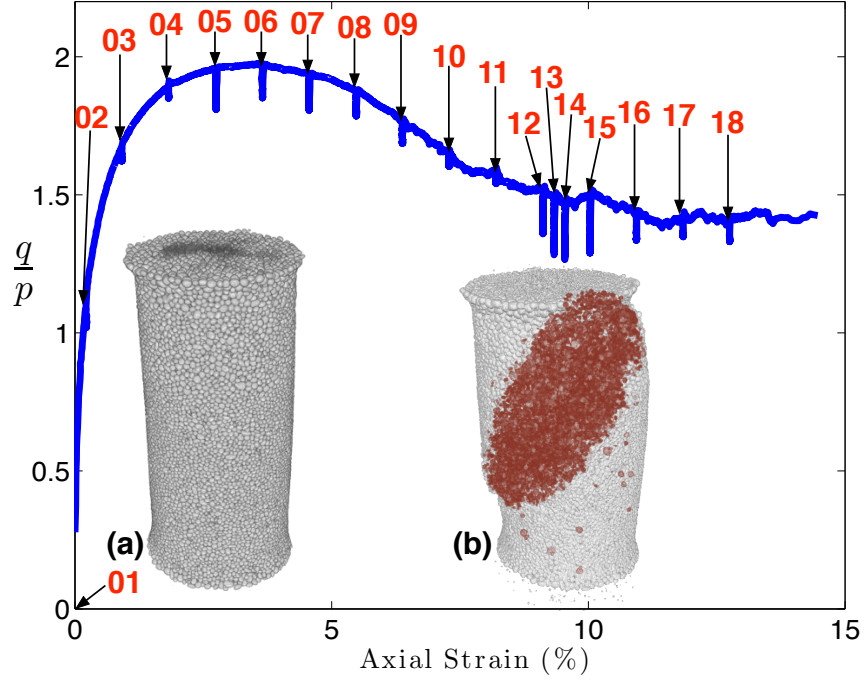


Figure 5.3: Macroscopic load-displacement response. Inserts: triaxial specimen at (a) initial state and (b) 9.11 % axial strain, with shear band highlighted. Station numbers are labeled in red.

mental displacements are observed in the elements at the specimen boundary, a product of the boundary effects stemming from the DIC technique itself and the specimen membrane; these edge elements are excluded from subsequent dilatancy calculations.

Following a finite element interpolation approach, we calculate the incremental strain over an element e as:

$$\Delta \epsilon^e = \mathbf{B}^e \Delta \mathbf{d}^e \quad (5.4)$$

where $\Delta \mathbf{d}^e$ is the incremental displacement vector containing the nodal incremental displacements (from experiment as described above) and \mathbf{B}^e is the strain-displacement matrix in Voigt notation. Incremental strains are calculated over each Gauss point (eight integration points per trilinear brick element) and then used to compute the average incremental strain over the element as:

$$\Delta \epsilon_v^e = \text{tr} \Delta \epsilon^e \quad (5.5)$$

$$\Delta \epsilon_s^e = \sqrt{\frac{2}{3}} \|\text{dev} \Delta \epsilon^e\| \quad (5.6)$$

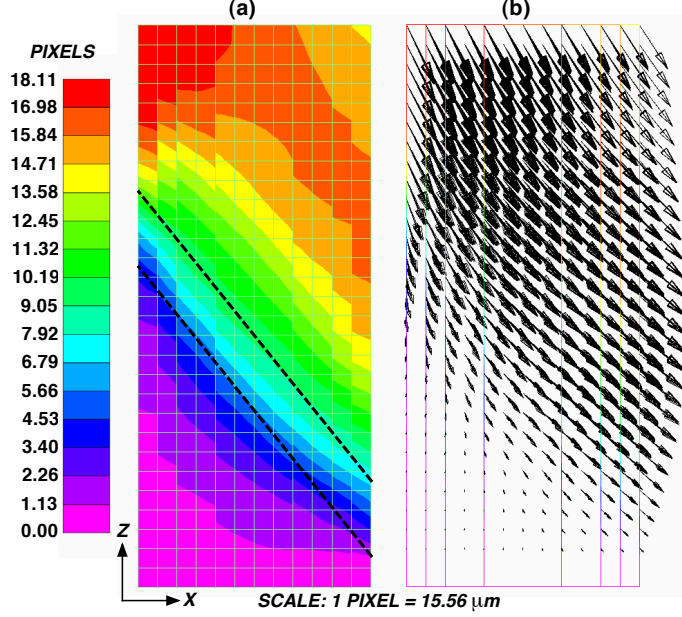


Figure 5.4: Incremental displacement field between tomographic stages 6 and 7 obtained from DIC: (a) contour of magnitude of incremental displacement vector with shear band within the region indicated by dashed lines and (b) incremental displacement arrows.

Based on the displacement field obtained from DIC, we identify the shear band region and all the elements within it. We note that the data shows that the deformations are fairly homogeneous throughout the shear band. With this, the average dilatancy in the shear band is calculated as:

$$\beta = \frac{\Delta \bar{\epsilon}_v}{\Delta \bar{\epsilon}_s} \quad (5.7)$$

and shown in Figure 5.5. Here, we have assumed that the elastic strain increments are negligible and their plastic counterparts dominate, allowing us to use the measured increments directly to quantify dilatancy.

5.3.3 Multiscale computation

To check that the average dilatancy evolution inferred from experiment is correct, an experimentally-driven multiscale computation was performed. Applying the hierarchical multiscale scheme proposed in [22], we upscaled the experimental average dilatancy evolution into a continuum finite element model of the triaxial specimen and the global structural responses of the model and experiment were compared.




Figure 5.5: Dilatancy inferred from DIC data with fit for multiscale calculation (see Section 5.3.3).

Underlying the multiscale scheme, we use a Drucker-Prager constitutive model with the yield surface and plastic potential being:

$$F(p, q, \mu) = q + \mu p = 0 \quad (5.8)$$

$$Q(p, q, \beta) = q + \beta p - \bar{c} = 0 \quad (5.9)$$

respectively, where p, q are the pressure and shear invariants of the stress tensor, μ is the generalized friction coefficient, β is the (plastic) dilatancy, and \bar{c} is a free parameter such that the plastic potential crosses the yield surface at the same stress state (p, q) . Generally, the microscopic variables μ and β are obtained either through an experiment or a discrete model. For simplicity and convenience, however, we can invoke the stress-dilatancy relation $\mu = \beta + \mu_{cv}$ and express β as a function of μ . The resulting multiscale scheme would then require three calibrated material parameters — the elastic modulus E , Poisson's ratio ν , and residual shear strength μ_{cv} — plus a dilatancy evolution β inferred from experiment. For the experiment considered here, the calibrated parameters are $E = 125$ MPa and $\nu = 0.3$, $\mu_{cv} = 0.81$. When these parameters and the shear band dilatancy evolution β in Figure 5.5 are upscaled into the finite element model, the global response of the experiment is reproduced as shown in Figure 5.6.

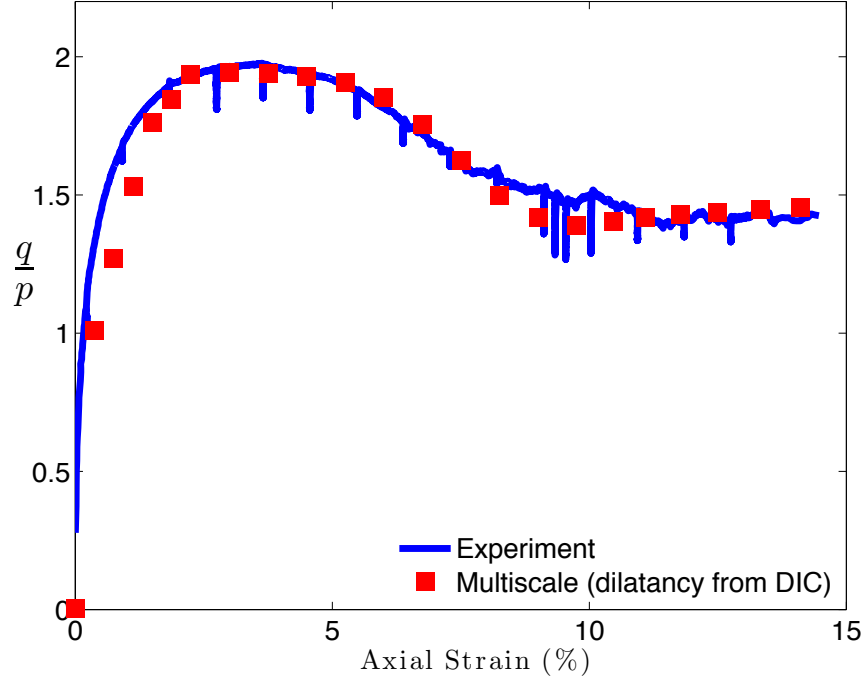


Figure 5.6: Global response from multiscale computation using dilatancy evolution information inferred from DIC (refer to Figure 5.5).

Remark 5.3.1 *The elastic parameters E and ν are determined based on the assumption that the material is linear-elastic and homogeneous up to about 0.09% global axial strain, as is apparent from the initial steep linear portion of the global stress-strain curve. We further assume that the dilatancy in the shear band is zero during the elastic stage and evolves as shown in Figure 5.5. The elements outside the shear band are assumed to evolve identically except that the dilatancy value is maintained once the peak is attained. This produces a state of inhomogeneous deformation after stage 6 (around 4% axial deformation), where the bulk of the deformation and global response of the specimen starts to be governed by the evolution of the shear band.*

Remark 5.3.2 *We note that the above experiment cannot be used to probe the stresses locally and only local deformation fields can be reconstructed from the experimental data. Hence, only dilatancy β can be inferred from the experimental data. As such, a comparison of the evolution of frictional resistance μ inside the shear band between experiment and a discrete model cannot be directly made. One can, however, determine the adequacy of a discrete model by checking if the model residual strength is close to this calibrated value of μ_{cv} (see Section 5.3.4).*

5.3.4 Discrete modeling

In the shear band identified in Section 5.3.2, two locations are arbitrarily selected for unit cell extraction. Starting from a 3D XRCT image at these locations, we apply the characterization step described in Section 5.2.1 to capture the morphologies of the grains. Subsequently, NURBS surfaces are fitted over the characterized grains for use in our NURBS-based DEM. Two discrete models are constructed (see Figure 5.7), each comprising of about 10^3 grains and with model thickness (~ 10 grains) that extends the full shear band thickness. We then apply the node generation procedure described in Section 5.2.2 with $\epsilon_d = 5$ voxels and $\epsilon_\theta = 10^\circ$, resulting in between 400 and 500 nodes per grain. For simplicity, we have constructed our model in voxel space. In this exercise, we have considered two unit cells due to limited computational resources.




Figure 5.7: Unit cells 1 and 2, generated through the characterization step using 3D XRCT data. The shaded grey region shows cell orientation in local coordinate system with respect to the global coordinate system.

The DIC data indicate that the homogeneous deformation of the shear band region is accompanied by dilation normal to the shear band plane. Therefore, we idealize the loading protocol for the discrete models (in the rotated coordinate system shown in Figure 5.8) as a plane strain shear with a vertical (along the 3-axis) confinement stress. It is known that dilatancy is affected by the confinement stress. For simplicity, however, the vertical stress was calibrated to be a constant. The assumption of a constant vertical stress does not affect the residual strength μ_{cv} since it is determined by particle morphology, and this has already been captured through the characterization process. Movement of grains on the model boundary in the shearing process is prescribed using two rotating smooth side walls, and a feedback loop is used to maintain a constant vertical stress on the top wall. Dynamic relaxation [57] is used where sufficient damping is introduced and the wall movements are

sufficiently slow to remain close to the quasi-static condition. Calibration of the discrete model parameters is performed on one of the cells (Unit Cell 1) while the second cell (Unit Cell 2) serves as a reproducibility check using the same calibrated parameters. The calibrated parameters are normal stiffness $k_n = 3 \times 10^6$, shear stiffness $k_s = 2k_n/3$, interparticle friction coefficient $\mu_p = \tan \phi_p = 0.18$, and vertical stress of 2 units (corresponding to 130 kPa). Due to the explicit time integration used in our NURBS-based DEM, we have made a trade-off between a high normal stiffness (which results in smaller time steps) and a shorter simulation turnaround time.




Figure 5.8: Deformed configurations for unit cells 1 and 2. Shaded grey region shows cell orientation in local coordinate system with respect to the global coordinate system: final applied shear angle is 40° from the 3-axis.

From the DIC data, a one-to-one correspondence between the axial global strain and applied shear angle can be established. The shear angle is approximately 40° (from the vertical 3-axis) corresponding to a global axial strain of about 12%. The average dilatancy evolutions calculated from the discrete models are compared with those inferred from DIC and as shown in Figure 5.9, the evolutions from both models match well with experiment. Also plotted in Figure 5.9 are the evolutions of stress ratio q/p , which reach approximately 0.8 at critical state and are close to the calibrated μ_{cv} value of 0.81 in the experimentally-driven multiscale calculation. Assuming that the two unit cells provide an adequate representation of the average dilatancy evolution and residual strength in the shear band, we can use this information and proceed with the same multiscale calculation described in Section 5.3.3. The calculation result (see Figure 5.10) shows that the predicted global structural response that uses discrete information is in excellent agreement with that which uses experimental data (refer to Figure 5.6). The above results point to the importance of and need for accurately capturing grain morphologies in order to improve the predictive capabilities of

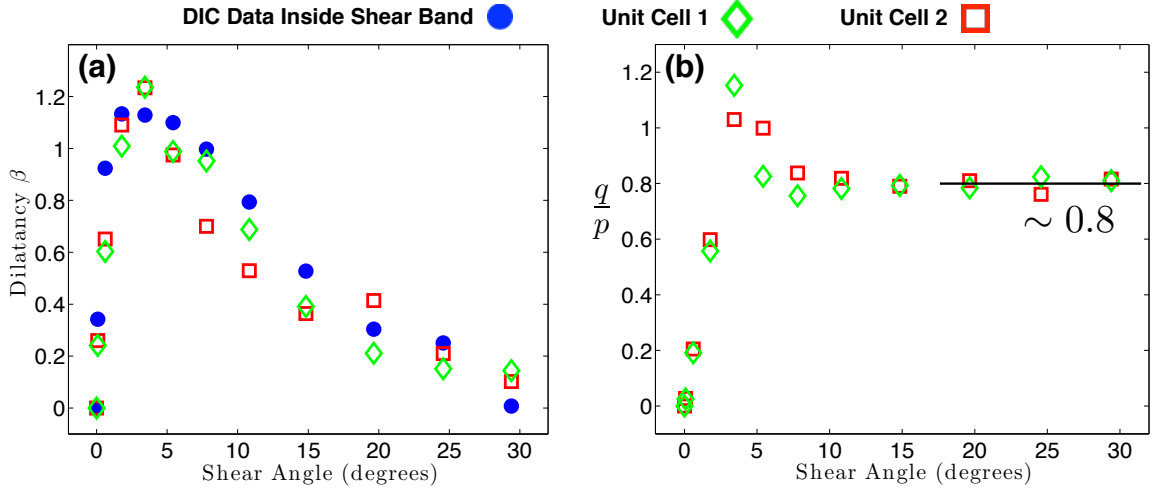


Figure 5.9: Results from discrete computations: evolution of (a) dilatancy and (b) stress ratio.

discrete simulation and multiscale techniques.

5.4 Closure

We have presented a computational mechanics avatar to probe the grain-scale behavior of granular materials. We have discussed the improvements in the characterization and discrete simulation components that have enabled us to transition from faithful representation of grain morphologies in X-ray tomograms of granular media to a morphologically accurate discrete computational model. We then applied and assessed the predictive capabilities of the avatar through a case study on a real experiment. We found good agreement between the microscopic quantities, namely the dilatancy evolution and residual strength, obtained from discrete simulations and those inferred from experiment. To our knowledge, this is the first quantitative comparison of microscopic quantities from discrete simulation and experiment in real sands. In addition, we found excellent agreement between the global continuum response calculated from multiscale computation using the extracted microscopic quantities and that measured in experiment. Overall, we found consistency in results across characterization, discrete analysis, and continuum response from multiscale calculations, providing the first and long sought-after quantitative breakthrough in grain-scale analysis of real granular materials.

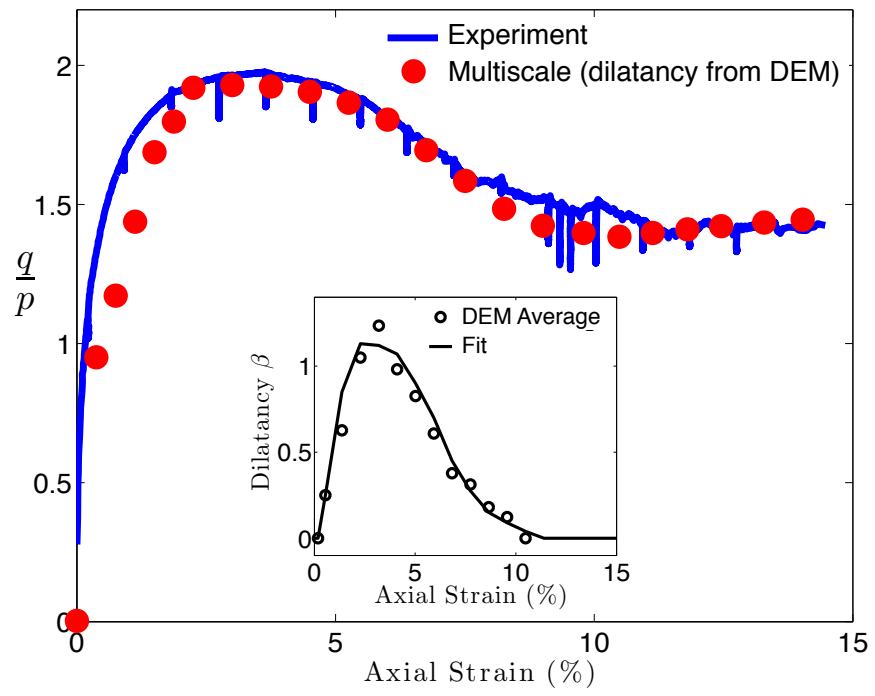


Figure 5.10: Global response from multiscale computation using dilatancy evolution information from discrete models (insert).