Chapter 6. Advanced Visualizations

This chapter reports the results of applying a novel image processing strategy (capable of identifying the three phases) to the raw reconstruction. The first section introduces the new processing strategy. The second presents the results in various forms that reveal the nature of trickle flow multiplicity. Thereafter, the void space is partitioned into distinct pores according to an established method. This allows the determination of pore-scale statistics which provide further insight into the nature of hydrodynamic multiplicity.

In reference to the overall strategy of this investigation, this chapter forms part of Part 2 (cluster-of-particles scale) of the Additional Experimental Insights block in Figure 1, and addresses Objective 2 from Chapter 3. At the end of this chapter, the list of characteristic multiplicity trends will be completed and the experimental investigation concluded.

6.1 Novel Image Processing Strategy[4]

6.1.1 Introduction

Note that in this discussion the objective is to obtain a high definition ternary image of the section of the bed under investigation where each voxel is associated with gas, liquid or solid. The new image processing strategy takes as its input the raw reconstruction, the origin of which is unimportant provided that the spatial resolution is sufficient to distinguish individual particles. Although cone beam geometry is used in this study, the proposed image processing techniques are equally applicable to tomographical

[4] The work in this section has been accepted for publication in Chemical Engineering Science, doi:10.1016/j.ces.2007.08.009, 2007.

reconstructions derived from fan or parallel beam geometries or magnetic resonance techniques.

The key toward improving upon thresholding is in realizing that it is a scalar operation where the assigned value of the voxel depends only on its intensity. It does not draw any information from the surrounding voxels. The alternative procedure developed here addresses this deficiency in two ways:

- First, by taking information from the surrounding environment into account in deciding whether the voxel is solid, liquid or gas.
- Second, by using the a priori knowledge of the particle shape in identifying solid voxels.

The methodology is to first identify the particle locations using all the available information (in three dimensions) from the drained bed. The inspiration for using all local information in deciding the value of a voxel is drawn from edge finding techniques in two-dimensions. Canny (1986) showed that faint edges in an image can be detected by recognizing that edge pixels usually appear along lines. A "faint" edge pixel is detected by using the "strong" edge information of the surrounding pixels.

The mathematical basis of the majority of steps in the proposed procedure is the (discrete) convolution operation, defined as:

$$C_{m,n,p} = \sum_{i=1}^{a1+b1-1} \sum_{j=1}^{a2+b2-1} \sum_{k=1}^{a3+b3-1} A_{i,j,k} B_{i-m+1,j-n+1,k-p+1} \qquad (17)$$

Here, entries into matrices are indicated by subscripts starting from 1. In computational terms, the $a1 \times a2 \times a3$ matrix $\boldsymbol{A}$ (in this case the reconstructed volume image) is acted upon by a $b1 \times b2 \times b3$ matrix $\boldsymbol{B}$ (called a convolution kernel) to yield a $a1+b1-1 \times a2+b2-1 \times a3+b3-1$ matrix $\boldsymbol{C}$, of which the innermost $a1 \times a2 \times a3$ matrix corresponds in position to $\boldsymbol{A}$ (i.e. $\boldsymbol{C}_{m,n,p}$ is located at the same physical position as $\boldsymbol{A}_{m-a1+1,n-a2+1,p-a3+1}$). One last aspect to consider is the values of $\boldsymbol{A}$ at the index values outside of $\boldsymbol{A}$ (i.e. $a1+1,$

a1+2, .. a1+b1-1 and similarly for the other dimensions and for ***B***). By equation 17, these values are required to compute the values of ***C*** at locations close to the boundary of matrix ***A***. These are assumed to be zero (i.e. ***A*** and ***B*** are zero-padded) which is appropriate since there is no image intensity outside of the bed. The convolution operation is written in shortened form as:

$$\boldsymbol{C} = \boldsymbol{A} \otimes \boldsymbol{B} \tag{18}$$

Here it is understood that ***C*** is resized (by neglecting its outermost shell) to correspond in size with ***A***. That is, entries into ***C*** therefore correspond in physical position to those in ***A***.

6.1.2 Image Processing Steps

The first objective is to distinguish the particles in an image without liquid flow. The processing of the CT image of a drained bed into a binary (solids-gas) image is done following the steps illustrated schematically in Figure 43 and discussed below. Figure 44 shows results of applying the steps to the image in Figure 40b (a two dimensional slice plane is shown although the procedure is applied in all three dimensions).

Filter

The first step toward identifying the solid in the tomography image is to apply a 3 × 3 × 3 median filter to the image. This is accomplished by specifying:

- ***I*** = original reconstructed intensity image (337 × 337 × 150 voxels)
- ***B*** = 3 × 3 × 3 matrix with 3^{-3} in each position (i.e. equal weight to each of the surrounding voxels)
- ***F*** = the filtered image (337 × 337 × 150 voxels)

F is simply the convolution of ***I*** with ***B*** (Figure 44a, to be compared to Figure 39). The purpose of the filter is to remove high intensity noise. Equally important, filtering smoothes the intensity transition from gas to solid. Smooth transition profiles are required for the next processing step. The filtering also removes some of the grainy quality to the

inter-particle void, and "flattens" the interior solid intensities (more uniform solid interiors reduce the number of "false edges" that register in the next processing step).

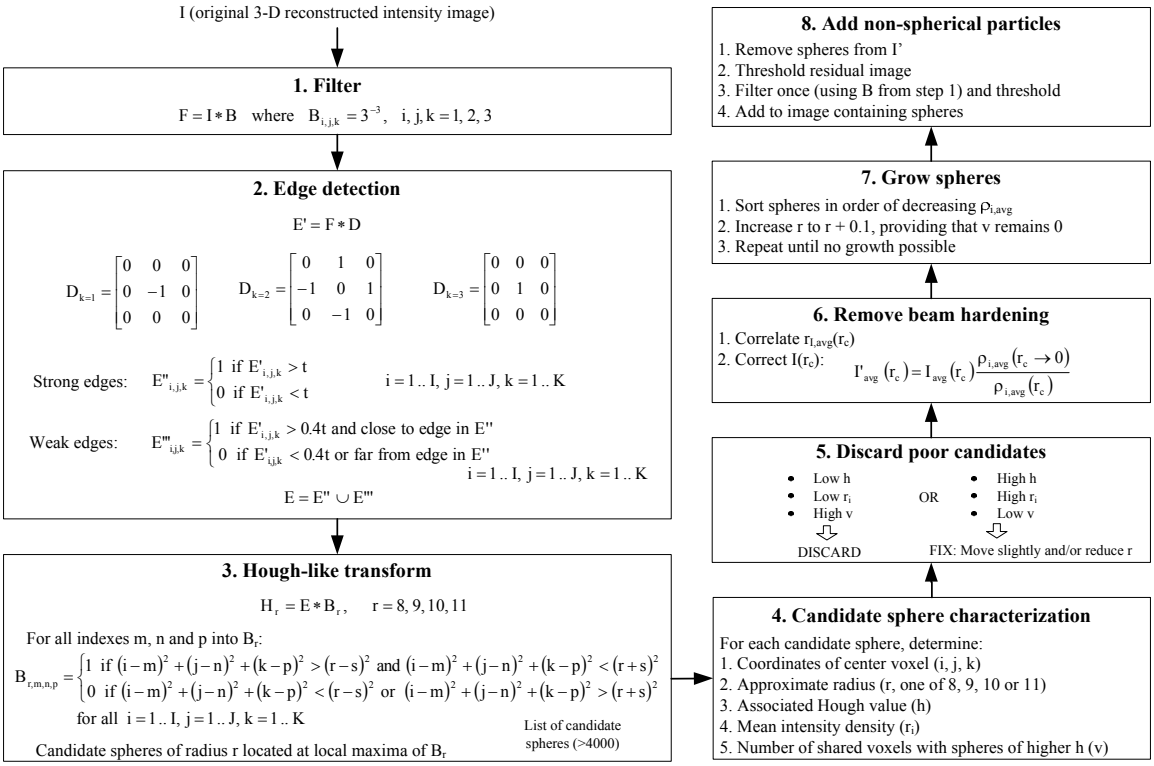

Figure 43. Image processing steps

Edge Detection

Edges are detected in images by identifying the locations of inflection in the intensity profiles (Van der Merwe et al., 2006), i.e. finding the values of *u* that satisfy:

$$\frac{\partial^2 \boldsymbol{I}}{\partial u^2} = 0\,,\ \text{where } u \text{ is the direction normal to the edge being detected.} \qquad (19)$$

The fact that the second derivative appears in this equation signifies the importance of the filtering step. Without pre-filtering the second derivative cannot be evaluated accurately due to the presence of high frequency noise. Ideally, the edges should be detected in directions normal to the particle surfaces, but the locations of the particles are not yet established.

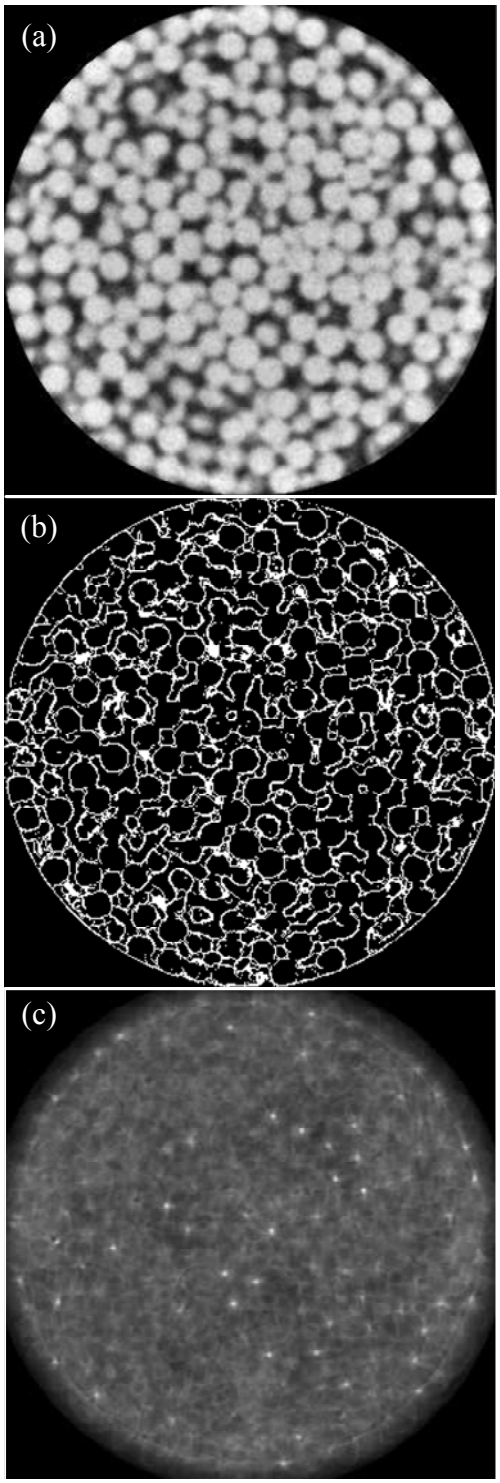

Figure 44. Cross-sections of a (a) Filtered image, (b) Edge image, (c) Hough-like transform

Therefore, the edges are detected in each of the three Cartesian directions that coincide with the image discretization (i.e. in each of the matrix directions). This is easily accomplished using convolution. A three-dimensional analogue of the Canny method (Canny, 1986) was found to yield superior results. It is a two-stage detection algorithm that finds first strong and then weak edges. A cross-section of the edge image ($\boldsymbol{E}$) is shown in Figure 44b. Large white patches in this image are instances where the slice plane coincides with an edge plane (e.g. near the top of a sphere).

Hough-like Transform

Examination of Figure 44b reveals that the edges of the spheres are neither smoothly nor explicitly accurately captured. Nevertheless, the (three dimensional) edge image can be used to identify likely coordinates in the volume that correspond to sphere centres. This is done by postulating a sphere radius (in this case 8, 9, 10 and 11 pixels were used; around 1.25 mm) and then "searching" for locations in the volume where edge voxels often appear approximately that distance from the location (the actual spheres are not exactly uniform in size so provision has to be made for spheres of different sizes). Computationally, this is accomplished by convolving the edge image ($\boldsymbol{E}$) with a convolution kernel ($\boldsymbol{B}_r$) that is constructed with 1's at a radius r from the centre of the image and zeros elsewhere. In practice, $\boldsymbol{B}_r$ is constructed with 1's whenever the distance from the central voxel is within $\pm s$ pixels of r. This is necessary because of two reasons. First, the digital cubic array (matrix) representation of a volume image means that spherical objects are inaccurately captured. Second, the particles can be expected to have small deviations from exact spherical shapes. In this way, $\boldsymbol{B}_r$ resembles a spherical "egg-shell" configuration with a "shell thickness" of $2s$ pixels (after some experimentation a value of $s = 1$ was adopted in this study). The original inaccurate edge image ($\boldsymbol{E}$) is convoluted with the "mask" of an accurate spherical shell yielding the Hough-like transformed image. The likely particle centres can be identified as the coordinates where the Hough-like image exhibits local maxima. The two-dimensional cross-section of the Hough-like image (with $r = 10$) is illustrated in Figure 44c. In this figure, the white spots

mark locations where one is likely to find a particle centre. With each Hough-like transform there is a radius associated, since the specification of $\boldsymbol{B}_r$ requires a postulated radius. A list of candidate spheres with their radii approximated to one of 8, 9, 10 or 11 pixels is therefore attained. There are more than 4000 such candidate spheres with normalized Hough-like transform values above 0.35 in the volume under consideration. A rough estimate of the real number of spheres in the volume can be obtained from the bed porosity (0.41, measured independently) and the approximate particle radius (1.18 mm):

$$N = \frac{V(1-\varepsilon)}{V_p} \approx 1545 \ particles \tag{20}$$

Many of the candidate spheres are computational artefacts or overlap with other candidate spheres. These need to be removed. An alternative to the Hough-like transform is the distance transform, where the convolution kernel of a solid sphere is used together with the filtered image $\boldsymbol{F}$. While the convolution of these two also yields maxima at potential sphere centres, the intensity profile through the centre has a slope approaching zero. In using the edged image, the Hough-like transform also identifies the sphere centres and the line profile slopes are very large at the centre. This means that the proposed procedure is more sensitive to the exact position of the sphere centre than the distance transform.

Candidate Sphere Characterization

The candidate spheres are characterized by:

1. Position and approximate radius.
2. Hough-like value (normalized) at centre of sphere (h, indicates the correlation with identified edges)
3. Mean intensity density (ρ_i)

4. Number of overlapping voxels (v) with candidate spheres of higher h (since candidate spheres with higher h values are more likely to correspond with the actual spheres).

Discard Candidate Spheres Unlikely to Correspond to Actual Spheres

Roughly speaking, the candidate spheres could be divided into two groups:

- Relatively low h, relatively low ρ_i or relatively high v
- Relatively high h, relatively high ρ_i and relatively low v

Sometimes phantom spheres are identified when there is a good correspondence with the edges of a combination of actual spheres. These spheres have low intensity densities and fall in the first group. In other cases, actual spheres are identified multiple times at one location but with different radii. Of these the larger ones (larger radius) generally have high v values and are removed on the basis that the spheres cannot physically overlap. Removal of the first group (low h, low ρ_i, high v) leaves a good approximation of the actual spheres in the bed. Some of these overlap ($v > 0$) with spheres of higher h values. In those cases where v is small ($v < 10$), visual inspection indicate that there is an actual sphere approximately in that position. In order to remove the overlapping, the sphere is moved locally (by no more than 2 voxel distances) to maximize the value of ρ_i. In all but a few cases, this reduced the value of v to 0. In the few remaining cases, the sphere radius was reduced by 1 voxel and then the value of ρ_i was maximized by translation. After these operations, 1526 spheres remained in the volume (which compares well with the approximate value of 1545).

Correct Beam Hardening

It can be observed that the average intensity density (ρ_i) of a sphere toward the centre of the column is lower than that of spheres at the wall. This restates the beam hardening effect and can be quantified as illustrated in Figure 45. The beam hardening effect is removed by adjusting image intensities as a function of the radial distance according to:

$$\boldsymbol{I}'_{avg}(r_c) = \boldsymbol{I}_{avg}(r_c)\frac{\rho_{i,avg}(r_c \to 0)}{\rho_{i,avg}(r_c)} \tag{21}$$

The original and adjusted images are compared in Figure 45. The beam hardening effect has been successfully removed.

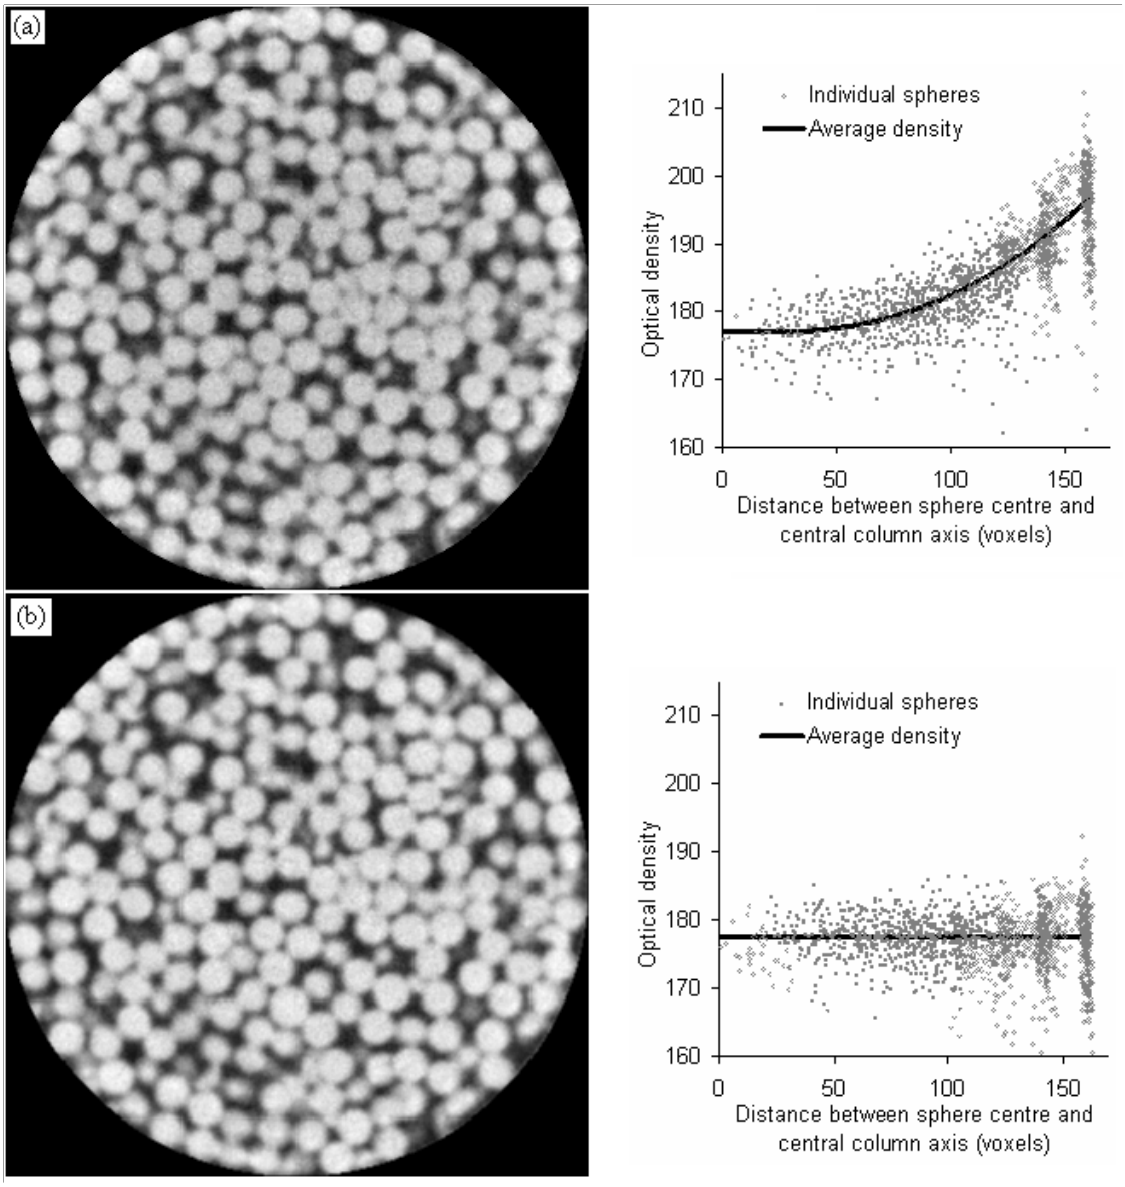

Figure 45. Beam hardening correction. (a) Uncorrected (b) Corrected.

Grow Spheres

The spheres identified thus far have approximate integer radii and do not overlap. In the previous steps, whenever two spheres with different radii at the same location were identified (by both having large Hough-like values), the smaller sphere was preferred. Inspection revealed that these spheres often do not touch any of their neighbours (a phenomenon which is again physically impossible since the spheres need to be supported against gravity). A simple way of correcting the radii is by growing the spheres in small increments until they touch. This is done by visiting a sphere, growing its radius by 0.1 voxels, evaluating whether the grown sphere overlaps with any existing spheres and only allowing the growth if it does not. The procedure is then applied to the next sphere in the list. Note that a single sphere cannot be grown until it touches the existing spheres, since this would not allow its neighbours to experience any growth. Note also that the order in which the growth process is applied determines which spheres are grown preferentially. It is necessary to appeal to the nature of computed tomography in order to determine which spheres to grow first. It is known that a sharp gas-solid transition reconstructs to a bi-directionally smoothed profile, schematically shown in Figure 46. In this one dimensional analogue, the mean intensity density of a candidate sphere of radius r is given by the area under the curve divided by r. If $r < r_{actual}$ (the actual radius), the intensity density will be high. Therefore, the candidate spheres with high ρ_i are those that should be preferentially grown. A cross-section of the resulting image that contains all the spherical particles is shown in Figure 47 (spherical particles shown in grey, ignore for the time being the white non-spherical particles).

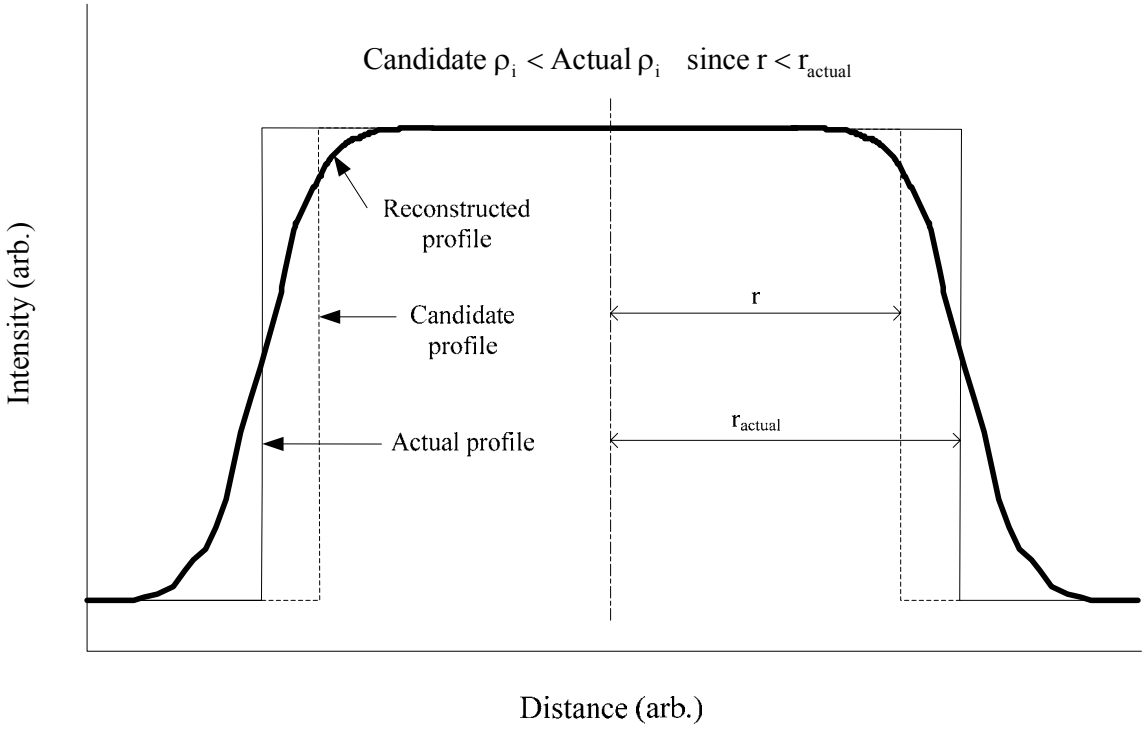

Figure 46. Motivation for the growth prioritizing

Procedure Generalization

The procedure stipulated above appears to correctly identify the locations and sizes of the spherical particles in the volume by using non-local information in establishing the nature of each voxel. A priori knowledge of the spherical shape of the particles was used in calculating the Hough-like transform (in particular in specifying $\boldsymbol{B}_r$). This step can be generalized to identify any particle of known shape by appropriate specification of $\boldsymbol{B}_r$. For example, if the bed had been packed with cylinders of different sizes, a Hough-like transform will be calculated for each size and orientation of the cylinders. This is equivalent to searching for cylinders of a certain size and a certain orientation in the bed (for each size and orientation a different $\boldsymbol{B}_{size,orientation}$ will be constructed). This expression is somewhat complicated for particles for which rotational symmetry does not hold. Nevertheless, although the procedure would be computationally expensive, it can be applied to any bed packed with particles of known shape and approximate size.

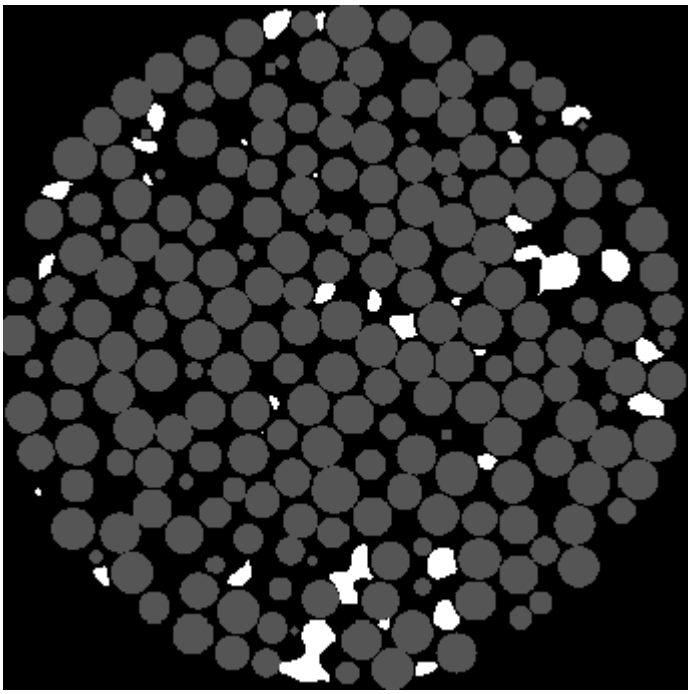

Figure 47. Solids image (non-sphericals shown in white)

Addition of Non-Spherical Particles

Although this procedure is effective for known shapes (in this case spheres), a significant fraction of the particles in the bed under consideration is not captured. Inspection reveals that these particles have arbitrary shapes (usually broken spheres). They are identified by taking all voxels with an intensity higher than some threshold, that is not already part of an identified sphere. While this is inaccurate there are relatively few non-spherical particles (as illustrated in Figure 47, non-spherical solid shown in white). In summary, the three dimensional volume image, $\boldsymbol{S}$, consists of a 3D matrix of voxels with a value of 1 if the voxel is located inside a particle and zero if it is not. In addition, the location and radius of each spherical particle is also known. This enables particle scale evaluations of the bed and flow structures. It is encouraging to note that the overall bed porosity

computed from this image (***S***) (0.41) compares well with that determined by gravimetric means (0.41).

Identifying the Fluid Phases

The liquid in the bed is isolated by subtracting the volume image of the bed under irrigation from the volume image of the bed without irrigation (Figure 48a). This image is then thresholded (at a value that removes an assumed Gaussian noise distribution) and the solids image (***S***) is superimposed onto it. Because the solid phase has been properly identified, the liquid located at particle contact points are now captured with greater accuracy. At some rare locations it was found that there is a slight gap (1 voxel in size) between the liquid phase and the solid. Based on visual examination of the original reconstructed images, these gaps are taken as computational inaccuracies and filled in with liquid. A cross-section of the final image (***T***) is shown in Figure 48b. The proposed procedure is recommended as a superior alternative to thresholding in CT image processing.

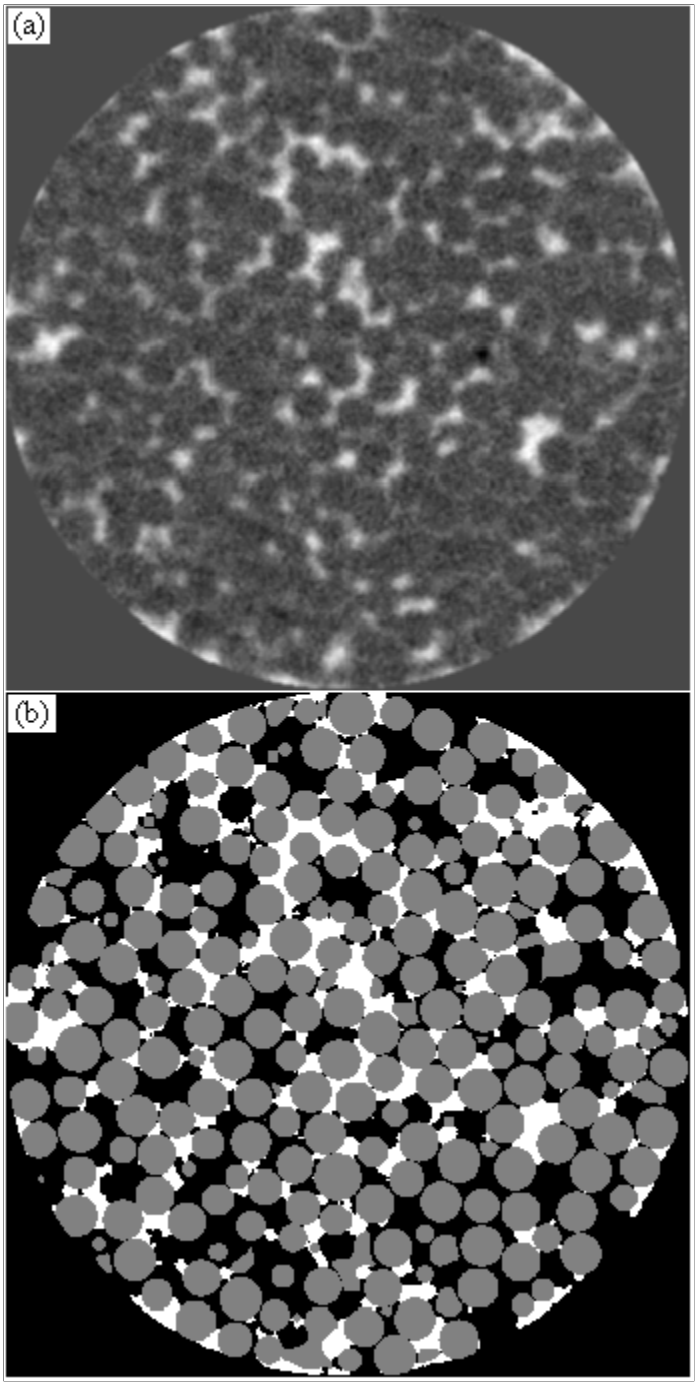

Figure 48. (a) Dynamic liquid in the bed obtained from subtracting the drained image from the flow image. (b) Ternary gated image (***T***) showing solids in grey and dynamic liquid in white.

De-Rectangularization

At this stage, the image (***T***) is a three dimensional matrix with each entry identified as either solid, liquid or gas. Most particles are approximately 10 voxels in radius. While volumes can be calculated in this representation by addition of voxel volumes, external surface areas cannot be calculated by addition of the areas of each external face. For example, a sphere with radius 10 voxels (i.e. 1.18 mm) has a real volume of 6.88 mm^3 and an external surface area of 17.5 mm^2. The sum of voxel volumes comes to 6.85 mm^3 (and approaches the real volume as the resolution becomes better), but the sum of external face areas is considerably higher than the real surface area (22.3 mm^2) and increases further as the resolution becomes better. That means that the external solid surface area as well as the liquid-solid surface area cannot be calculated by counting faces. Instead, the surface area is calculated as follows:

- The edge voxels are identified as voxels that share at least one face with gas or liquid (wet edge voxels share at least one face with a liquid voxel). Each edge voxel has associated coordinates (*ex, ey, ez*) taken at the centre of the voxel. From standard geometrical arguments, the rotational and azimuthal angels of that voxel (θ, ϕ) can be calculated. Note that the distance from the centre of the central voxel is not necessarily equal to the radius, *r* (it is less than one voxel size shorter).
- The solid external surface is then "paved" with squares (patches), with the centre of the patch located at (θ, ϕ, *r*). Larger spheres therefore comprise more of these patches because there are more edge voxels. The areas of these patches are approximately equal ($1.81 \pm 0.02 \times 10^{-8}$ m^2) because of the regular spacing of the edge voxels.
- The specific external surface area calculated from these patches (1461 m^2/m^3), compares well with the same area calculated using the porosity and the average particle radius (1446 m^2/m^3). The gas-liquid interface is handled in the same way.

- Whether or not an external solid surface patch is contacted by liquid is determined by matching the (θ, ϕ) values of the wet edge voxels to those of the patches. The fraction of these patches contacted by liquid is the wetting efficiency.

There is one apparent error in this method. This lies in determining whether an edge voxel is wet or dry. Edge voxels close to the dry-wet interface on the particle surface will very often have some of its external faces in contact with liquid and others contacted by gas. In these cases, the voxel has been counted as both wet and dry. An estimate of the error in doing this can be made by counting these voxels as either wet or dry and then calculating the change in gas-solid or liquid-solid surface area. These changes were typically below 3% and the error is therefore negligible. A similar situation arises for liquid voxels that contact both the solid and the gas. Because it is smaller than the other areas, the error in the gas-liquid area is estimated to be around 12%.

It is important to note at this stage that the wetting efficiency calculated in this way is the dynamic wetting efficiency (because the drained bed that was used as a subtractor includes the residual liquid holdup). It was found to be very difficult to distinguish the residual (sometimes referred to as static) liquid in the difference between a drained bed and a dry bed. This is because the pendular rings are very small for this system and because the internal holdup in the drained bed changes the magnitude of the edge effect discussed earlier (the operation of subtracting a dry bed from drained bed therefore yields a lot of noise at the particle contact points). Therefore, the residual liquid holdup (RLH) was added to $\boldsymbol{T}$ by an alternative approximate method. It is known that residual holdup accumulates at locations of close particle-particle contact, and that for this particular system the majority of RLH is in the form of pendular rings at the contact points. The RLH was measured independently for this system at 2.5% by the weighing method. The RLH image ($\boldsymbol{SL}$) is added to $\boldsymbol{T}$, by calculating it as $\boldsymbol{SL} = 1$ where $\left(\left[\boldsymbol{S} * \boldsymbol{F}\right] - \boldsymbol{S}\right) > t$. This yields artificial RLH at locations where the local solid density is high (e.g. contact

points). The filtering and thresholding parameters are set so that the average RLH equals the 2.5% previously measured. The final image showing the solid, gas, residual liquid and free-draining liquid is shown in Figure 49. This is a considerable improvement over Figure 41, and can now form the basis of a comprehensive investigation in trickle flow behaviour. A great deal of information is available from the three dimensional image of the bed where each phase has been more accurately identified, and is reported in the next section and explored further in Chapter 7.

6.1.5 Image Processing Strategy Conclusions

Computed tomography techniques offer unprecedented potential for detailed examinations of trickle flow hydrodynamics. This section showed how more advanced image processing techniques can be employed to extract valuable information from less than perfect tomography results. X-ray CT results of a pilot scale trickle bed under various different conditions are used to illustrate the procedure. Using the proposed approach, it is possible to investigate trickle flow hydrodynamics to a level of detail that has been impossible before.

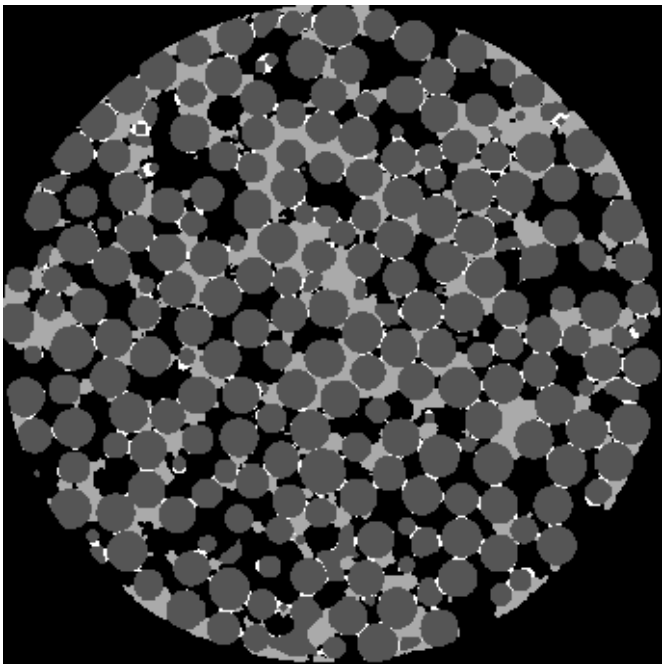

Figure 49. Addition of RLH (in white)

6.2 Tomography Results

It should be obvious now that a wealth of information can be extracted from the high detail tomographic images. This includes bed characteristics like porosity and its axial and radial distributions, particle coordination numbers, the fluid flow patterns, global and local holdup and its distribution, average and particle scale wetting efficiency and its relationship to bed and flow characteristics, gas-liquid interfacial area and the influence of velocity changes, pre-wetting and flow histories on the hydrodynamics. Some examples follow. These are all pertinent to the discussions in Chapter 7 and are intended to illustrate both the power of the experimental technique and the nature of hydrodynamic multiplicity.

Figure 50 shows cross-sections of the bed for three of the modes. Notice the differences in liquid distribution. Figure 51a shows the results of imaging the bed in three dimensions without any liquid present (4.7 mm of the reconstructed 18 mm are shown in the axial direction). If the liquid is imaged without the presence of the solid, the prevailing flow patterns become apparent (Figure 51b to Figure 51h). A careful examination of these figures reveals the following:

- There are un-irrigated volumes in the Levec mode at low liquid velocity, and these volumes become active at high velocity. The Kan-Liquid and Super modes show very good liquid distribution at both high and low velocities, but smaller scale un-irrigated volumes persist in the Levec mode at high velocity. There appear to be little difference in macroscopic liquid distribution between the Super and Kan-Liquid modes.
- An increase in velocity results in higher saturation and increased spreading of the liquid – a greater fraction of the bed volume is occupied by liquid.

Images like those in Figure 51 serve to illustrate global features, but it is necessary to augment them with quantitative analyses. From an image like that depicted in Figure 51a, the radial and axial porosity distributions can be calculated. They are shown in Figure 52 along with the trend predicted by the recent model of De Klerk (2003b) for similar systems. The radial porosity distribution was evaluated at different length scales, but no new information that relates to hydrodynamic multiplicity was gained from such analyses.

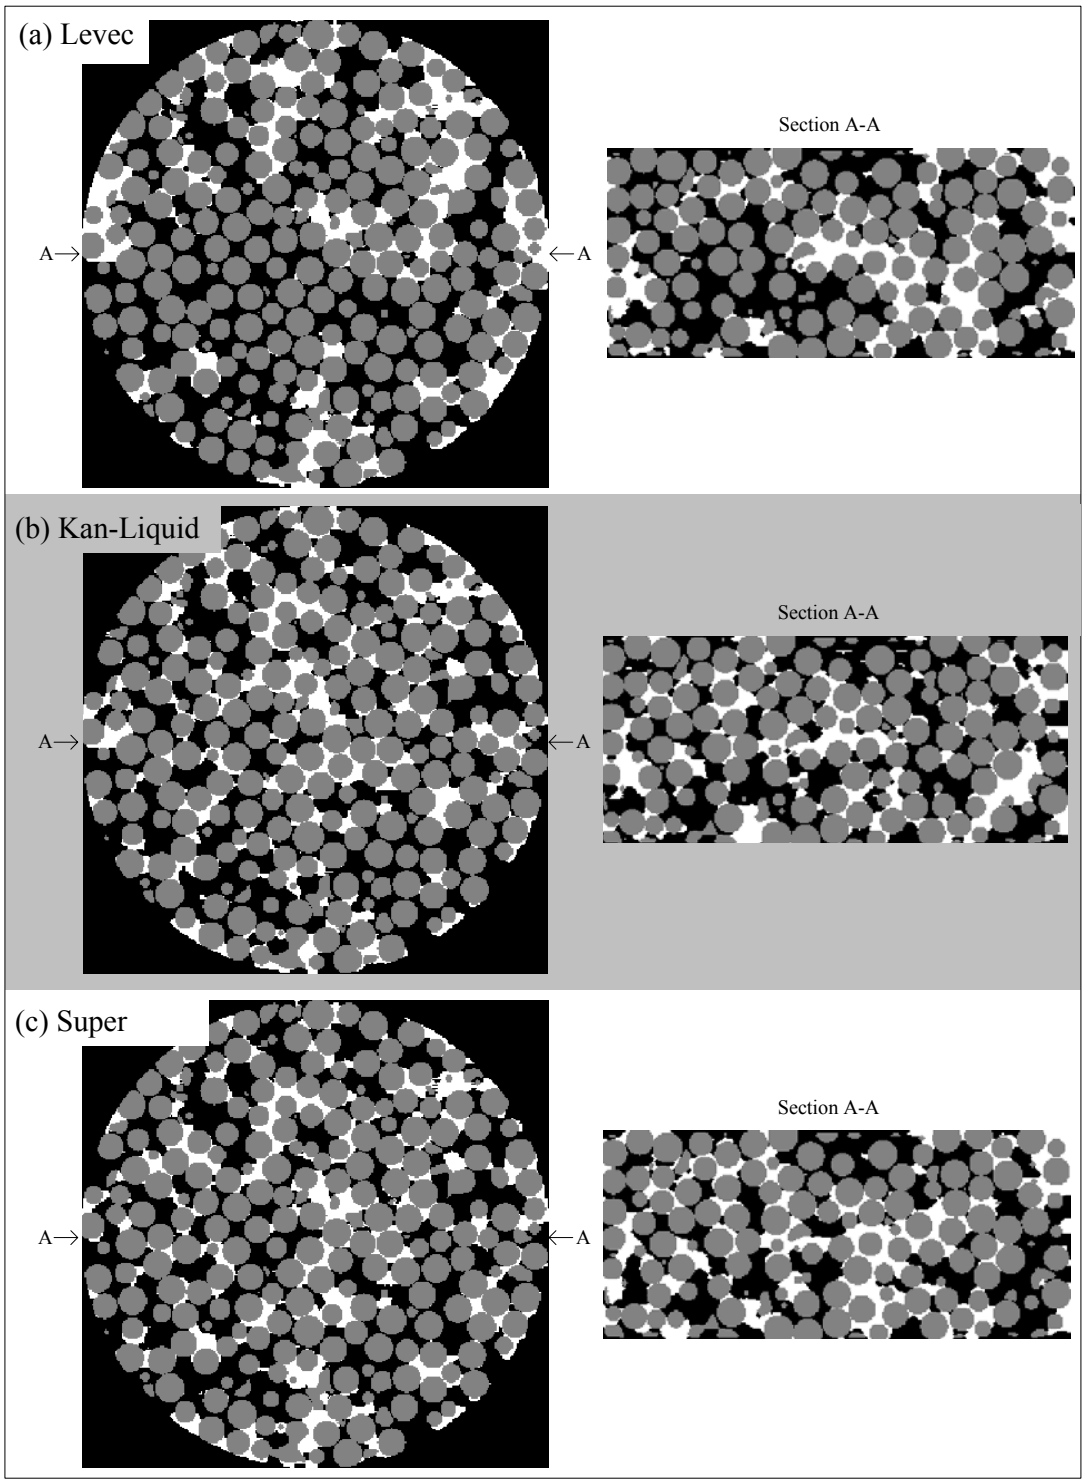

Figure 50. Tomography results at u_L = 2.7 mm/s and u_G = 4.7 cm/s

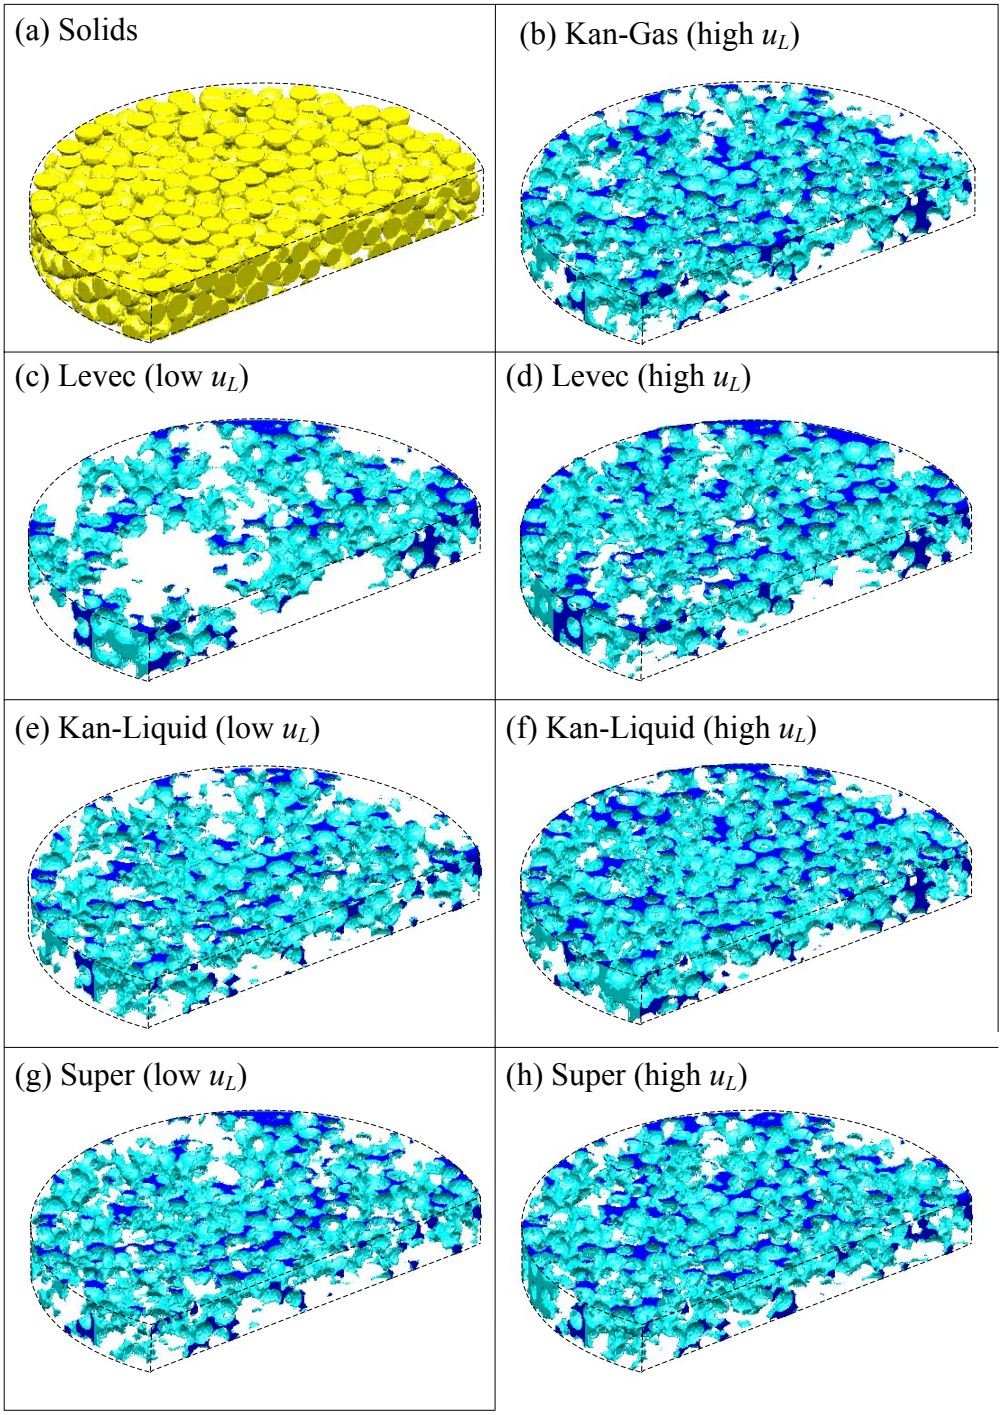

Figure 51. Three dimensional images of the (a) solid (yellow), and (b-h) the liquid (surface in light blue, slice plane in dark blue) in the bed in the different modes. High u_L = 5.3 mm/s, low u_L = 2.6 mm/s, u_G = 3.2 cm/s.

The axial porosity distribution is sinusoidal (not damped) with a relative standard deviation of 3.3%.

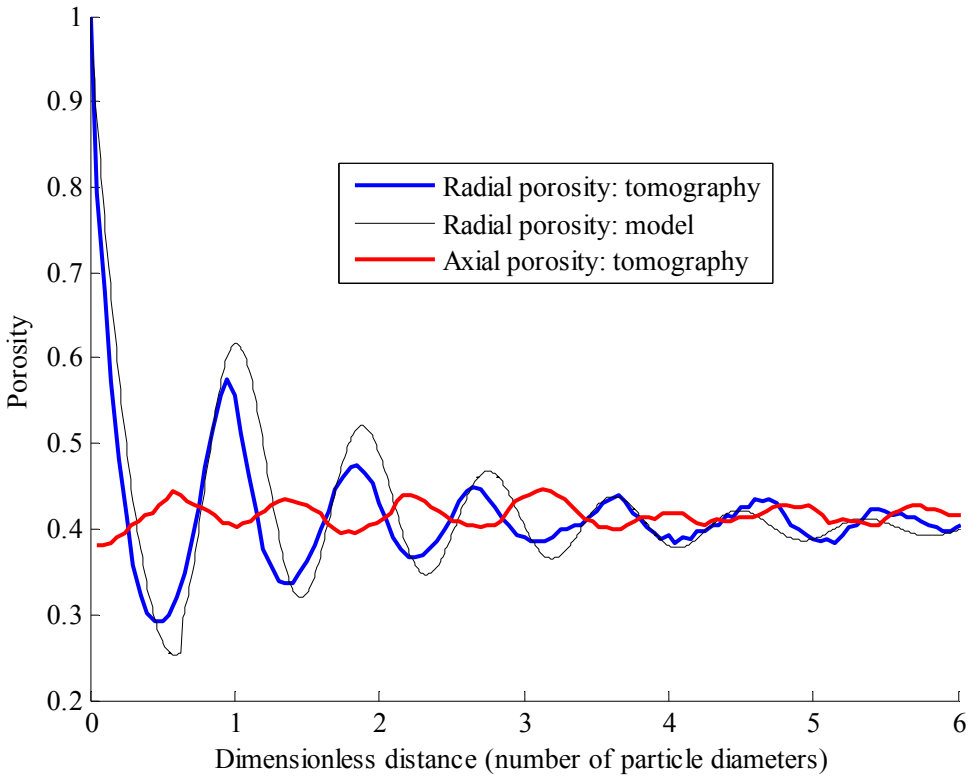

Figure 52. Porosity distribution in the radial and axial directions.

Other bed-scale properties that can be extracted from the three-phase images are the liquid holdup, wetting efficiency and gas-liquid mass transfer area (Figure 53). To keep the figures un-cluttered, only the Levec and Kan-Liquid modes are shown here (the Super and Kan-Gas modes data are available but is not shown). The data generally agree very well with the trends evaluated in Chapter 4. This instils the belief that the images are indeed representative of the hydrodynamics. The level of detail inherent in the processed tomographic images allows the determination of localized hydrodynamic phenomena, an example being the wetting geometry on individual spheres (Figure 54). However interesting these kind of visualizations are, the focus is now shifted specifically toward establishing what the mechanism of hydrodynamic multiplicity is, as well how it explains

the qualitative behaviour that has been observed so far. This issue is dealt with in context in Chapter 7.

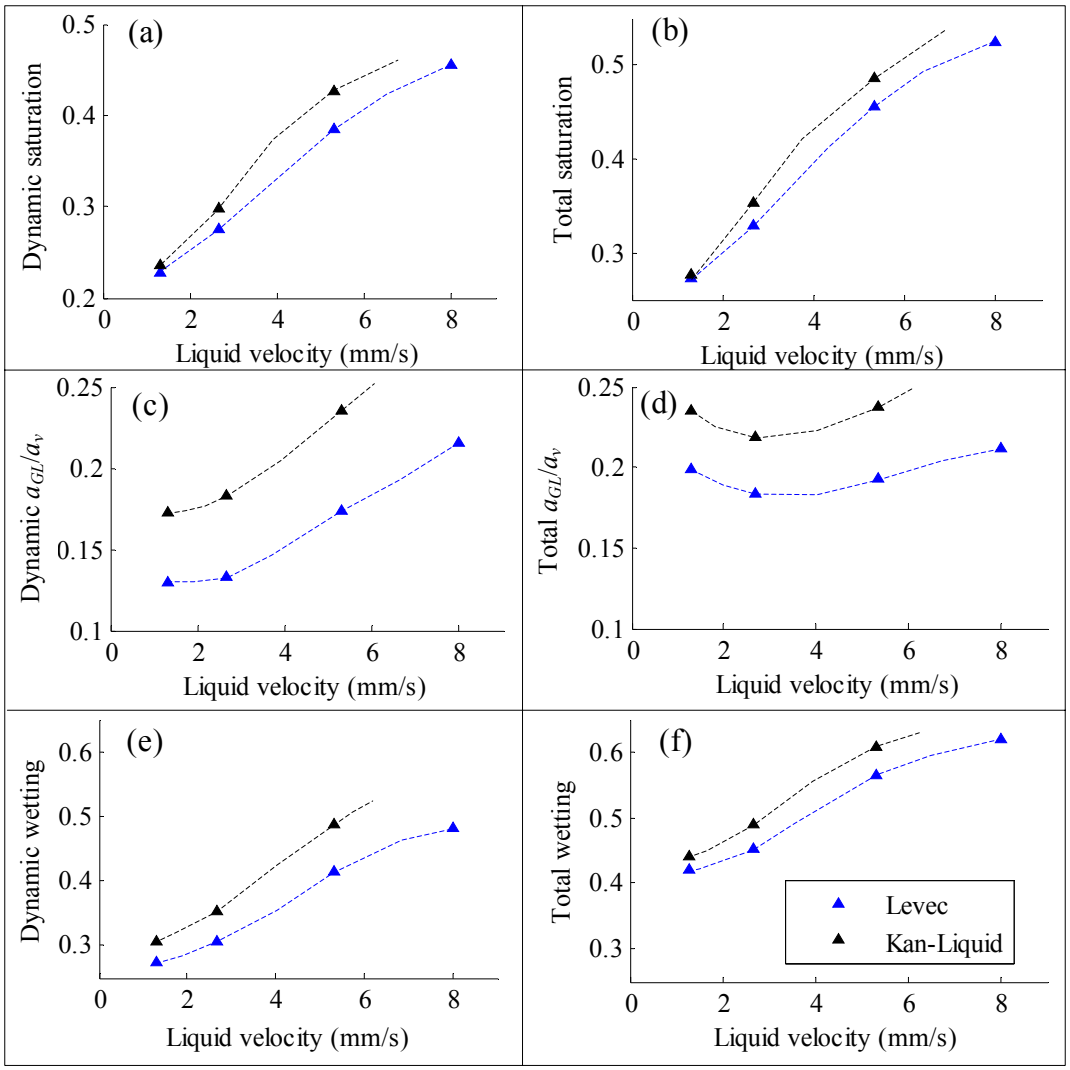

Figure 53. Bed-averaged hydrodynamic parameters computed by the method developed in this study (u_G = 1.6 cm/s). (a), (c), (e) is the dynamic saturation, gas-liquid area and wetting efficiency, (b), (d), (f) is the total saturation, gas-liquid area and wetting efficiency.

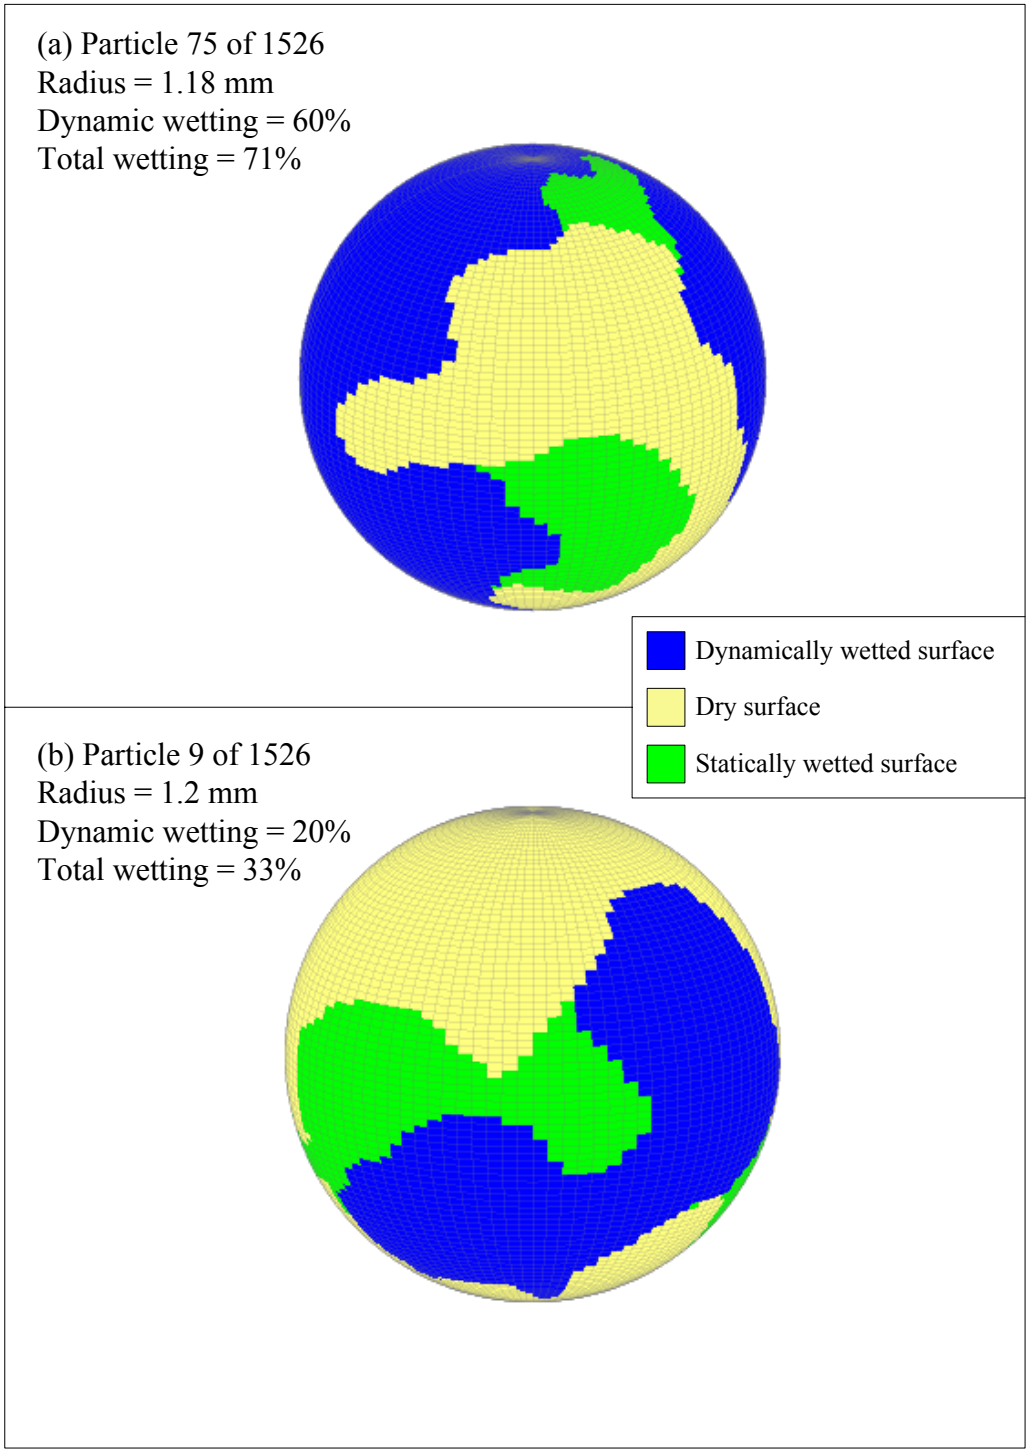

Figure 54. Three-dimensional images of typical partially wetted spheres (dry surfaces in yellow, dynamically wetted surface in blue, statically wetted surface in green).

6.3 Bed Structure and Flow Pattern Characterization

Additional information can be extracted from the high definition ternary image generated in section 6.1. This is done by partitioning the bed into "pores" according to the procedure developed by Baldwin et al. (1996). The details of this procedure are discussed in detail in literature and are only summarized here briefly. Pores are defined as relatively large void volumes bounded by solid surfaces and planes erected where the hydraulic radius is a minimum (Dullien, 1992). Starting from the three dimensional image where each voxel is identified as either gas or solid (***S***, no flow), the void space is "thinned" morphologically. This operation is basically the discrete version of the distance transform and yields high intensities at locations far from any solid (i.e. the potential centres of the pores). The pores are then grown from these local maxima sequentially, i.e. the largest pore centres (as identified by the value of the thinned image at the local maxima) are grown by one voxel in each direction outward from the centre, then the largest pore and the next largest is grown by one voxel and so on until the whole void space is filled. A pore stops growing when it reaches the solid surface or meets another pore. Two touching pores therefore always share an approximate plane of voxels (which is referred to as a pore neck or pore exit). These steps are illustrated in Figure 55 for a 2D section of the actual bed. A technical detail is that in each case the "connectivity" of voxels have to be chosen for the growth step. The three options are 6-connected (voxel faces touch), 12 connected (faces and edges touch) or 26-connected (faces, edges and corners touch). It was found that the resulting pore classification did not differ drastically with these options and 26-connection was used. Baldwin et al. (1996) states that this procedure results in an unambiguous pore classification, but there is some ambiguity in what exactly is meant by a "local maximum" in defining pore centres. As an example, Figure 56b shows a surface plot of the thinned image for the region shown in Figure 56a (the same region as used before). In this figure, the decision to view point 2 as a local maximum results in a smaller pore centred at point 1 (Figure 56c).

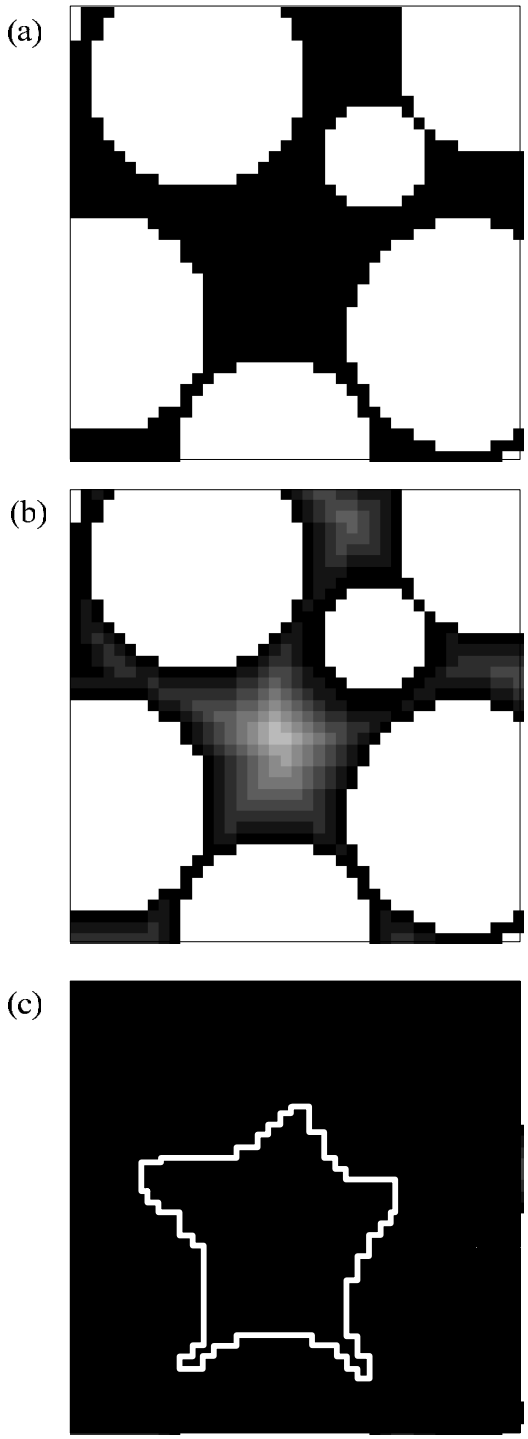

Figure 55. Pore space classification (a) Original image (b) Thinned image (c) Pore (white line shows pore boundary)

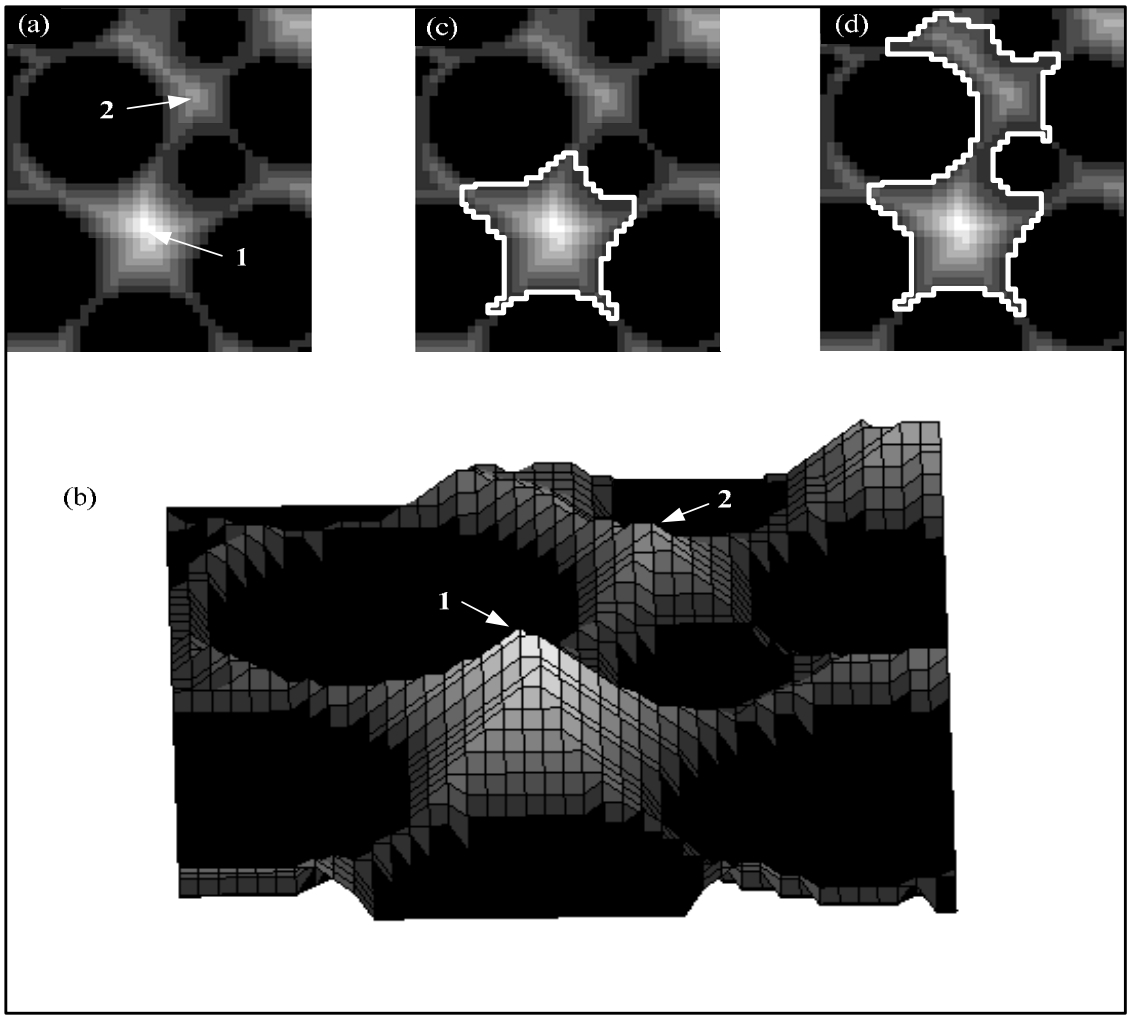

Figure 56. The ambiguity of pore classification. (a) Thinned image (b) Surface plot showing local maxima (labelled 1 and 2) (c) Pore classification with small window (b) Pore classification with large window.

If point 2 is not considered a maximum, the pore centred at point 1 is much larger and encompasses the region around point 2 (Figure 56d). One has to choose a window size that is to be considered in order to identify local maxima. If the window size is small, both points 1 and 2 will register as maxima. If it is large, only point 1 will register. For the remainder of this chapter, a window size of 10 voxels (i.e. one particle radius) is chosen as this yielded results that seemed intuitively correct. However, all calculations were repeated with smaller window sizes and while the data are obviously different, none of the conclusions drawn regarding hydrodynamic multiplicity were compromised.

An example of a three dimensional pore is shown in Figure 57 in relation to its neighbourhood. The figure also shows the liquid inside the pore. From this sort of analysis, a number of pore-scale statistics regarding the bed and the liquid distribution in the bed can be calculated. These pore-scale statistics include the following (as shown in Figure 58):

- The centroid of each pore
- The volume of each pore - the pore volume distribution
- The external surface area of each pore, including the pore-solid area and the pore-pore area. This allows the calculation of a pore hydraulic diameter (6 times the pore volume divided by the pore-solid area), which shows an approximate normal distribution (average = 1.8 mm, relative standard deviation = 14%).
- The coordination number of each pore (the number of pores each pore is connected to). Also included here is a list of which pores are connected and the characteristics (area, location and circumference) of each pore-pore interface (neck).
- The sphericity of each pore (while the particles are spherical, the pores are far from it).

When the liquid distribution in each pore is also considered, some additional statistics can be calculated (as shown in Figure 59):

- The holdup distribution (by pore) and its evolution with gas and liquid velocity in each multiplicity mode.
- The wetting distribution (by pore) and its evolution with gas and liquid velocity in each multiplicity mode.
- The gas-liquid area distribution (by pore) and its evolution with gas and liquid velocity in each multiplicity mode.

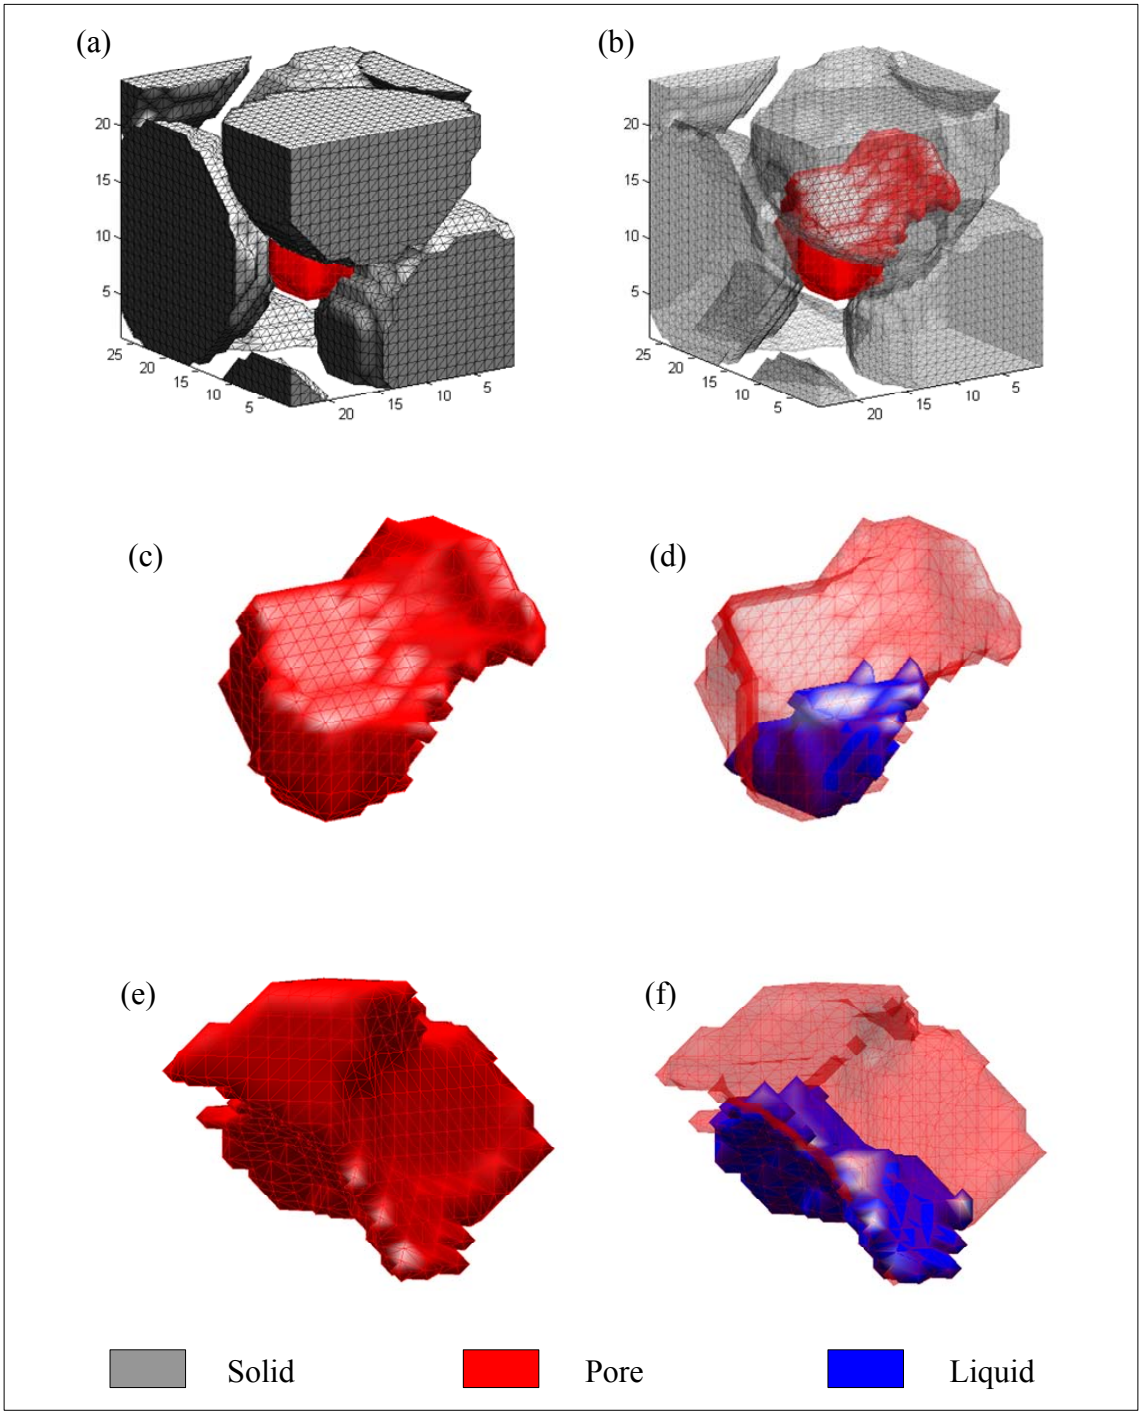

Figure 57. An example of a typical pore, (a) in relation to its neighbourhood (b) the pore nestled in the void space (c) the pore on its own (d) the liquid inside the pore (e) the pore from another angle (d) the liquid inside the pore from another angle.

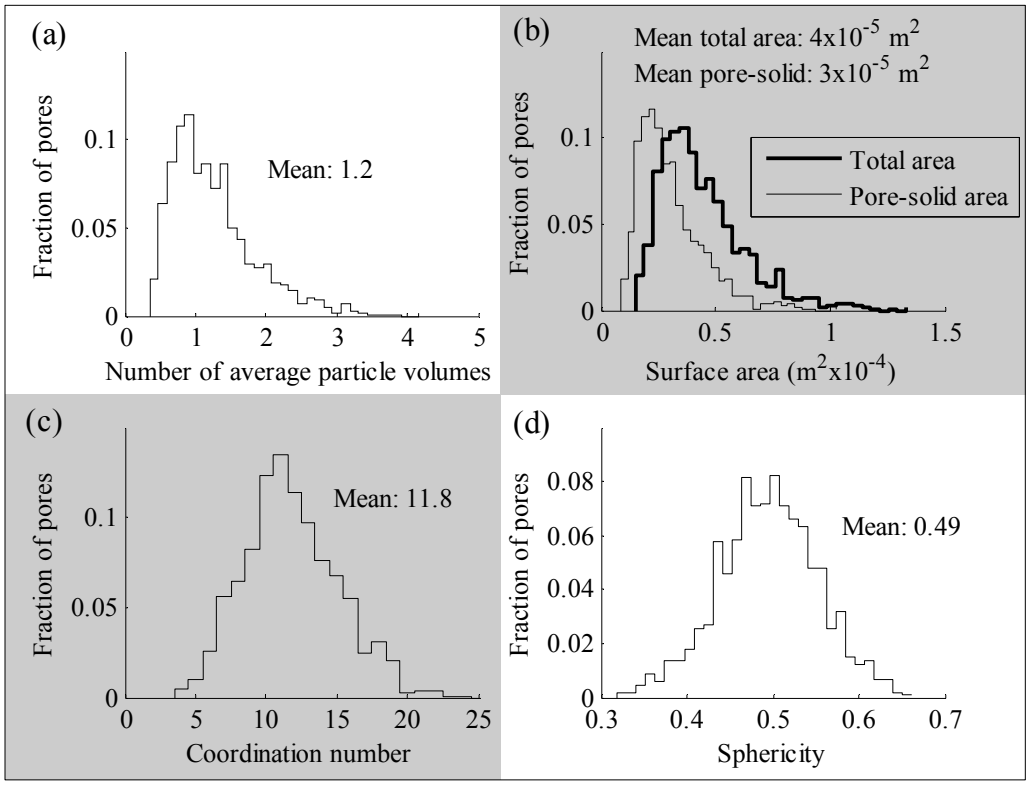

Figure 58. Pore-scale statistics of the packed bed. (a) Pore size distribution (b) Pore external surface area distributions (c) Pore coordination number distribution (d) Pore sphericity distribution.

Similar trends to those observed in Figure 59 were observed for the other gas velocities and they are not shown for the sake of brevity. The pore holdup distributions are of particular importance because they are indicative of the mechanism of liquid distribution (as will be discussed shortly). Two additional multiplicity trends are:

(S) The Levec mode holdup distribution is always bi-modal (and increasingly so as liquid velocity is increased). The Kan-Liquid mode is uni-modal at low velocity and bi-modal at the highest velocity. See Figure 59.

(T) The fraction of pores that contain no dynamic liquid in the Levec mode decreases from 12.7% to 5.5% and then to 0.2% as velocity rises from 1.3 to 2.7 to 5.3 mm/s. In the Kan-Liquid mode, this fraction is always approximately zero – even at low velocity (see Figure 60). Note also that increasing the gas velocity in the Levec mode results in a decrease in the fraction un-irrigated pores, despite also resulting in a decrease in holdup (as discussed in Chapter 4).

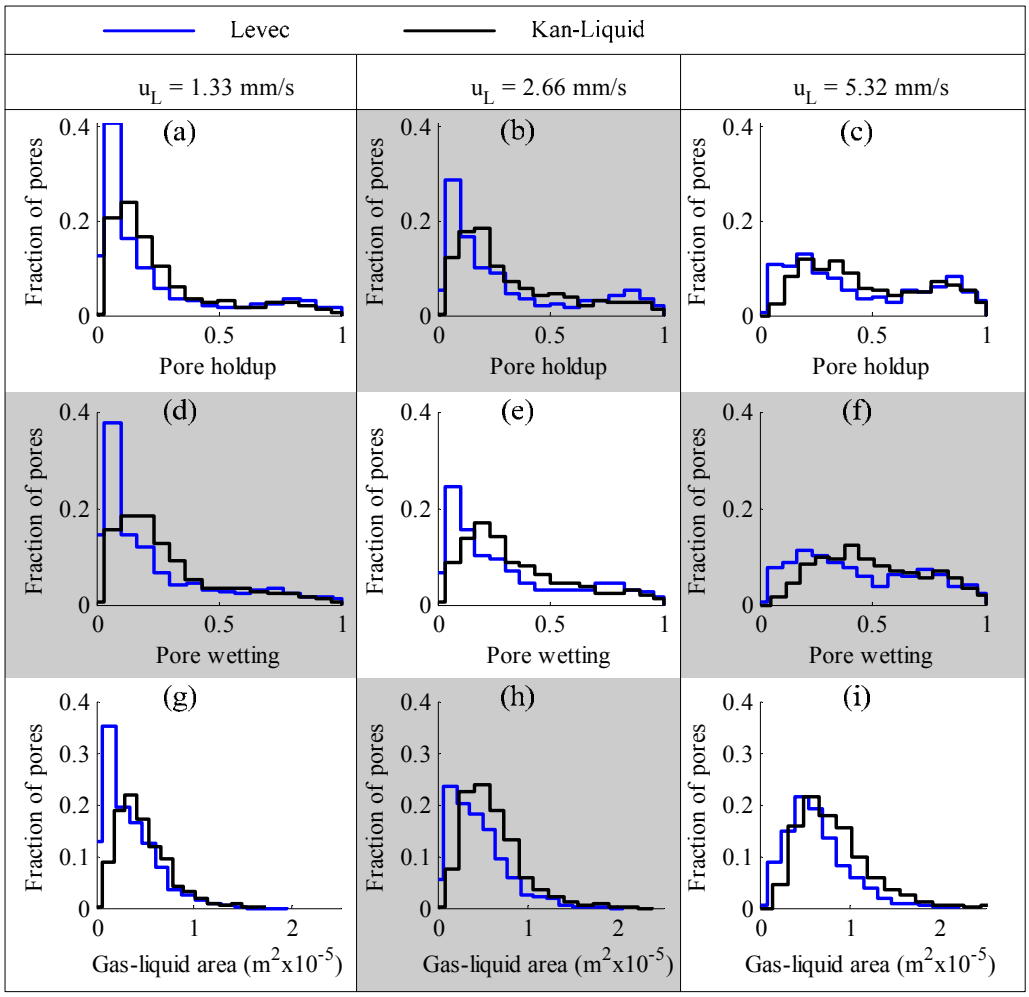

Figure 59. Pore-scale statistics of the flow pattern in the Levec (blue) and Kan-Liquid (black) modes. (a-c) Pore holdup distribution with increasing liquid velocity, (d-f) Pore wetting fraction distributions, i.e. the fraction of the solid surface exposed to the pore that is dynamically wetted, (g-i) Pore gas-liquid area distribution. u_G = 4.7 cm/s for all figures.

The trends listed in **(A)** to **(T)** in Chapter 2 to 6 are the fundamental features of hydrodynamic multiplicity as evaluated by the limiting case framework introduced in this work.

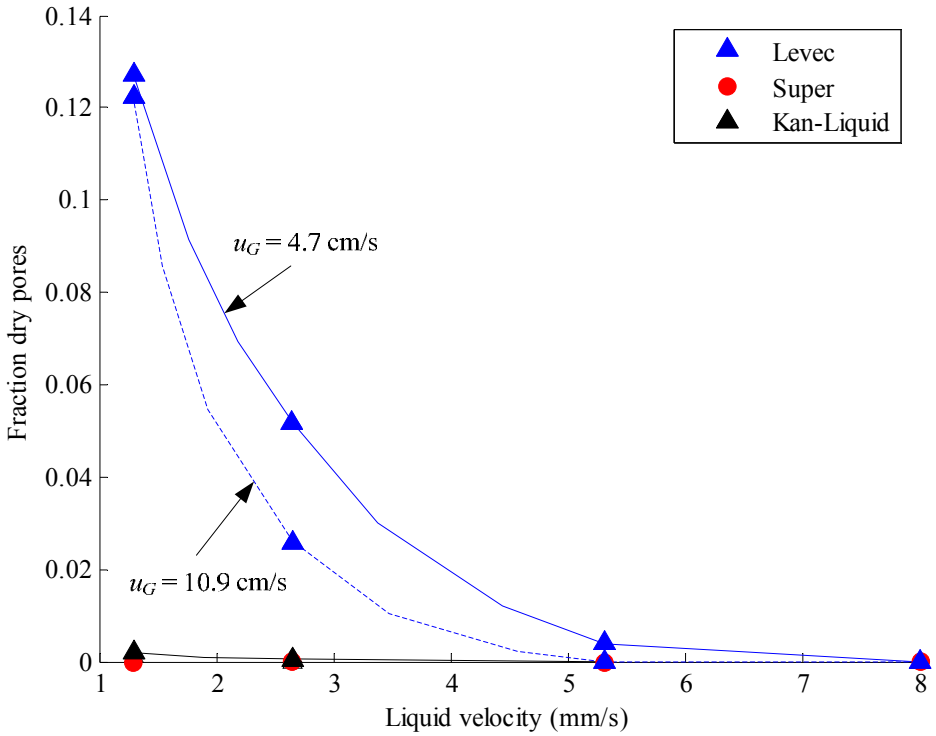

Figure 60. Fraction of pores without dynamic holdup in them (dry pores) as a function of gas and liquid velocity and hydrodynamic mode.