Resolution-invariant measurement of small objects in polychromatic CT data

Richard A. Ketcham*[a]
[a] Department of Geological Sciences, Jackson School of Geosciences, University of Texas, Austin, TX, 78712-1722

ABSTRACT

Features such as particles, pores, or cracks are challenging to measure accurately in CT data when they are small relative to the data resolution, characterized as a point-spread function (PSF). These challenges are particularly acute when paired with segmentation, as the PSF distributes some of the signal from a voxel among neighboring ones; effectively dispersing some of the signal from a given object to a region outside of it. Any feature of interest with one or more dimensions on the order of the PSF will be impacted by this effect, and measurements based on global thresholds necessarily fail. Measurements of the same features should be consistent across different instruments and data resolutions. The PVB (partial volume and blurring) method successfully compensates by quantifying features that are small in all three dimensions based on their attenuation anomaly. By calibrating the CT number of the phase of interest (in this case, gold) it is possible to accurately measure particles down to <6 voxels in data acquired on two instruments, 14 years apart, despite severe artifacts. Altogether, the PVB method is accurate, reproducible, resolution-invariant, and objective; it is also notable for its favorable error structure. The principal challenge is the need for representative effective CT numbers, which reflect not only the features of interest themselves, but also the X-ray spectrum, the size, shape and composition of the enclosing sample, and processing details such as beam-hardening correction. Empirical calibration is the most effective approach.

Keywords: 3D measurement, point spread function, shape characterization, resolution-invariant measurement, effective attenuation coefficients, segmentation, gold, economic geology

1. INTRODUCTION

In many applications of CT-based analysis, it is desirable to inspect and measure small features in larger contexts. For example, when imaging pore networks with the aim of modeling fluid flow [1], one may wish to image a volume several cm wide to obtain a representative volume, but quantify flow through pores spaces and throats on the order of 1's to 10's of μm. Similarly, quantitative analysis of trabecular bone [2] requires capturing a large enough volume to be mechanically representative while retaining fine struts and plates that, though small, can still affect elastic response. However, when feature size becomes small with respect to the resolution of the data, defined not only by the voxel size but also by blurring across voxels, measurements become more difficult and accuracy harder to obtain. This blurring, which can be summarized as a point-spread function (PSF), or essentially a smoothing kernel, is an inevitable feature of CT data that arises from a range of factors, such as focal spot size and drift, detector size and crosstalk, gantry imprecision, and the reconstruction process [3].

The effects of limited resolution on CT measurement of fractures have long been recognized, and a number of methods have been proposed to compensate [4-8]. These methods use the continuum information in CT numbers to estimate an aperture based on the CT number deficit compared to solid material, and allow apertures to be measured down to a fraction of a voxel width. An analogous set of methods have been developed in the medical field for measuring bone that is thin compared to the data resolution [9-11]. These applications are somewhat straightforward in that they are making an essentially 1D measurement (aperture, thickness) traversing a feature that is small in one dimension but large in the other two. The underlying principle of both set of methods is conservation of attenuation: the total attenuation anomaly for the feature is contained in the CT data, only dispersed. Although only approximately correct in polychromatic CT data due to beam hardening, it is close enough to improve measurements greatly.

* ketcham@jsg.utexas.edu; phone 1 512 471-6942; fax 1 512 471-9425; www.ctlab.geo.utexas.edu

A recent contribution [12] addresses the analogous problem of objects that are small in all three dimensions, using a 3D extension of the same principles. The case it considers is gold particles, which are economic to mine at ppm or g/t concentrations, and are typically sparsely disseminated in their host rocks [13]. To locate and observe a good sampling of gold particles in their natural context (e.g., associations with other mineral phases, veins, and fractures) often requires imaging a volume considerably larger than the particles themselves. However, even in this situation it remains desirable to accurately measure particle sizes and shapes to grade the ore and plan optimal extraction and processing methods [14].

This paper summarizes and discusses the Partial Volume and Blurring (PVB) method introduced by [12], and compares it against more the typical measurement technique based on thresholds. Global thresholds are inherently flawed for segmenting features that have one or more dimensions at or smaller than the scale of the PSF, as the threshold that will correctly recover the volume of such a feature is dependent upon its size and shape.

2. METHOD

2.1 Point-spread function

A number of methods are available for characterizing and quantifying the PSF of CT data [3]. This work uses the convention by [8] based on the formula:

$$PSF(x,y,z) = \exp[-8((x-x_0)^2 + (y-y_0)^2 + (z-z_0)^2)/r_{PSF}^2] \quad (1)$$

in which the PSF is defined as a discrete normal (Gaussian) curve with a characteristic dimension r_{PSF}, which is equivalent to 4-sigma, and (x_0, y_0, z_0) are the voxel coordinates. This definition of r_{PSF} enables a number of useful intuitions: 95% of the signal from a voxel is contained within a sphere of diameter r_{PSF}, and 95% of the transition between two adjacent materials along a flat interface occurs over the distance r_{PSF}, allowing it to be estimated visually. It can be more precisely calculated from the edge response function, using the method outlined by [15]. Although the PVB method calculations do not employ the PSF directly, knowledge of r_{PSF} provides guidance on data processing and interpretation.

2.2 Segmentation

Segmenting a feature of interest in the PVB method requires capturing all voxels affected by the feature. Thus, rather than trying to delineate the feature boundary with the surrounding matrix, as with typical segmentation, the region to be captured is expanded by a degree that depends on the PSF: either r_{PSF} /2 beyond the feature's true boundary, or r_{PSF} beyond the central region where the feature achieves its endmember (unblurred) CT number. This can be done using the "threshold and expand" technique [16]. An initial threshold is set at a CT number unique to the phase of interest, and selected regions are then expanded by a finite number of voxels in all directions to a second threshold just above matrix values (or below, if dark features such as pores are being measured). Additional measures to account for or encompass bright and dark streaks from metal artifacts are not needed or desirable, as their effect is minor and their origin is detached from the "conservation of attenuation" principle.

2.3 Conversion to volume

For each segmented region associated with a feature of interest, the PVB-corrected volume (v_{PVB}) is determined by summing the CT signal in the captured voxels, according to:

$$v_{PVB} = v_{vox} \sum p_{f,i}; \quad p_{f,i} = (CT_i - CT_{mat})/(CT_f - CT_{mat}) \quad (2)$$

where v_{vox} is the volume of a voxel, $p_{f,i}$ is the proportion of the signal in selected voxel i attributed to the feature of interest, CT_i is the CT number of voxel i, CT_{mat} is the CT number of the surrounding matrix, and CT_f is the CT number of the feature of interest.

2.4 Calibration

Equation (2) requires two calibration parameters, the CT numbers for matrix and the feature of interest. CT_{mat} is usually straightforward to determine, using for example a region of interest (ROI) encompassing a large number of voxels representative of the matrix surrounding the features of interest. If the matrix is strongly inhomogeneous, a local value can be calculated for each particle by averaging the CT numbers in the annulus immediately surrounding its segmented region. CT_f can be more problematic, and a number of approaches are possible. Possibilities include using CT numbers in the centers of the largest features of interest, provided they are large enough to not be blurred with surrounding material;

calibration based on reproducing the volume of a "known" particle; or first-principles calculation based on X-ray attenuation. For the study outlined here, we use the known-particle method and discuss the alternatives.

2.5 Re-thresholding

Once the PVB-corrected volume for a feature has been determined, its shape and orientation can be quantified by estimating its discrete extent using a customized, effective threshold. The number of voxels to capture is determined as $N_f = v_{PVE}/v_{vox}$, and then the threshold is found by sorting the selected voxels by CT number and using the N_f-highest value (or lowest, when processing negative features) as the effective threshold.

3. DATA

3.1 Gold particles

A series of twelve natural gold grains from the Witwatersrand (South Africa) with various sizes and shapes were weighed and 2D-imaged under a light microscope. Particle masses ranged from 12.5 to 334.3 µg, and maximum dimensions from 168 to 729 µm. Minimum dimensions could not be determined directly due to the irregularity of particle shapes, but estimates based on mass, estimated density, and projected area range from 42 to 119 µm. Density was estimated based on an assumed grain purity of 90% gold and 10% silver, as is typical from that locality. The gold particles were placed in a cylinder with quartz powder and cemented using epoxy infiltration to make a 10-mm-diameter core. Some iron filings, probably from the crushing plates used for quartz powder preparation, are also included in the core. Further details on this sample are provided in [12].

3.2 CT data

The core was imaged in two sets of scans, a series of five in 2004 using a BIR ACTIS scanner [17], and a series of seven in 2018 using a Zeiss Xradia MicroXCT. The 2004 scans were acquired at 180 kV, with the core isolated and embedded in cylinders of quartz powder with progressively increasing diameter from 20-50 mm, to approximate the effects of the grains being embedded in a progressively larger core. The 2018 scans were acquired at 80 kV and 150 kV, with the core isolated and embedded in a 35-mm-diameter core with quartz or garnet powder, to examine the effect of different host rock properties. The 2018 scans were reconstructed using beam-hardening corrections (linearization; [18]) to level the background and reduce streak artifacts. Voxel sizes ranged from 10-50 µm. The ACTIS scans had measured r_{PSF} values averaging 5-5.5 voxels, while the MicroXCT r_{PSF} averaged 3.5 voxels, indicating sharper images. Figure 1 shows an example slice and 3D rendering. Further information on scanning conditions is provided in [12].

Figure 1. (left) Example CT image of artificial quartz + epoxy core with embedded gold and iron particles. Gold grains are accompanied by moderate to severe streaking artifacts. (center) Re-windowed version of left image, showing gold grains closer to true dimensions. Overall variations in brightness are due to variable degrees of blurring based on particle size and shape, though the flat grain in bottom center and enlarged in inset has a dimmer center and brighter ends due to beam hardening. (right) 3D rendering of core. Quartz is semi-transparent white, gold particles are yellow, iron shavings are light blue, and air bubbles in epoxy appear as darker regions.

3.3 Data processing

The twelve scans were processed in Blob3D software [16], which has been augmented to include PSF-measuring and PVB calculations. As described in Section 2.2, gold grains were segmented by selecting a CT number sufficiently high to

encompass all gold grains without capturing any iron particles, and expanded by 2 voxels to capture the surrounding blur region, terminating at CT values just above the quartz+epoxy matrix. CT_{mat} was determined by averaging a large region of the matrix devoid of particles and bubbles. The effective CT number for gold was determined for each scan using one particle (grain 11) as a known, which allows determination of the CT_f value that will reproduce its volume.

To provide the most favorable and equivalent basis to compare PVB-method results with what would be obtained from a global threshold, a threshold was selected for each scan to reproduce the volume of the same known grain as used for the PVB calculation.

4. RESULTS

4.1 Volume

Figure 2 shows the results for particle volume. For data from each scanner (ACTIS, MicroXCT), scans are arranged from left to right in order of increasing voxel size (10-20-30-40-50 µm ACTIS; 15-15-35-35-35-50 µm MicroXCT), and thus decreasing resolution. Grain 11 is the "known" particle, and thus all measurements for it are equivalent.

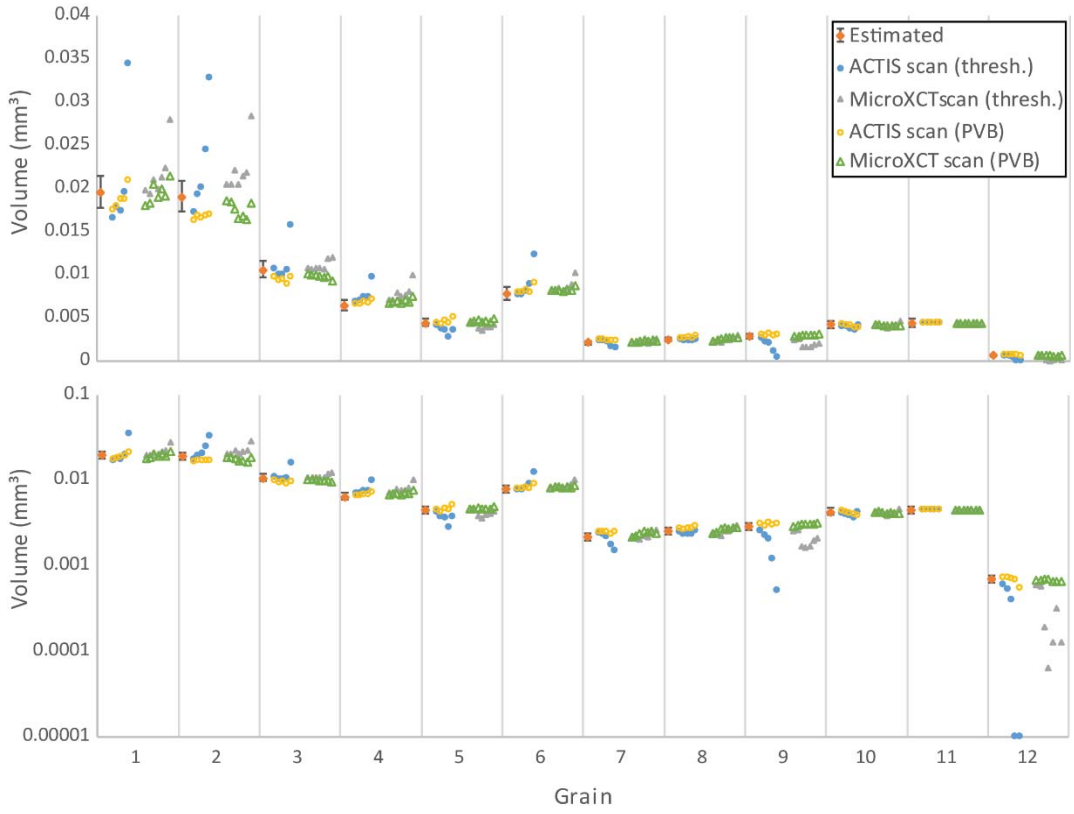

Figure 2. Linear and log plots of volume measurements based on the PVB method versus a global threshold. Estimated volumes are based on measured mass and estimated density, assuming a composition of 90% Au and 10% Ag.

The PVB-method measurements (open yellow and green symbols) are reasonably consistent across all scans and resolutions, showing only limited variation for the largest particles. The volume of particle 12 is approximately 5.5 voxels in the lowest-resolution data set for both scanners; its volume is recovered well in the MicroXCT scan but slightly worse in the ACTIS data, probably due to the latter's larger r_{PSF}.

The threshold-based measurements (solid blue and gray symbols) show a number of expected behaviors. As resolution falls, the sizes of the largest grains rise, indicating that the grain-11-based threshold was too low for them. This effect diminishes as particle size becomes similar to grain 11. Conversely, as particles get smaller than grain 11, decreasing resolution results in diminishing measured volume, indicating that the grain-11-based threshold is too high for those grains.

The two lowest-resolution ACTIS scans for particle 12 omit it entirely. Overall, resolution effects are more severe for the ACTIS scans, due to their higher r_{PSF}. Similarly-sized grains sometimes behave differently due to shape effects; the less compact a grain is, the more it is subject to resolution effects. For example, grain 9 has a branching morphology that distinguishes it from the other grains with comparable volume, increasing the relative impact of blurring.

4.2 Effective threshold

The effective thresholds for particles from several scans are shown in Figure 3, plotted as the ratio of their sphere-equivalent radii to the r_{PSF} for their respective scans, versus the fractional CT number difference between matrix and gold. Ideally, the appropriate CT number for a threshold is the midpoint between the object of interest and its surroundings [3], or a normalized value of 0.5. For these scans, the appropriate threshold to recover volume is just over 0.4 for the largest particles in the highest-resolution scans, and falls to below 0.1 for the smallest particle in the lowest-resolution scan. Much of the scatter in the data trend is due to shape effects. Typical of natural gold, most of the particles are flat and irregular, similar to wrinkled and torn sheets or toroids, while some feature irregular branching.

Also shown on Figure 3 are model curves resulting from convolving a 3D PSF with a number of shapes (cube, sphere, square plate, and hollow cubic and spherical half-domes) at a range of relative scales and finding the threshold that reproduces their volumes. Overall, the data follow the predicted trends fairly well, although none of the model shapes is capable of representing the complexity of the gold grain shapes fully. In general, shape-based effects will be increased by high surface to volume ratios, and the presence of sharp corners or protrusions that are more likely to get rounded due to blurring.

Additional shape considerations, and the effect of the PSF on recovering accurate particle shape and orientation information, are discussed in [12]. In general, shape metrics based on caliper-type measurements are more resilient to resolution effects than ones based on perimeters and surfaces areas. Shape anisotropy determinations generally require that the shortest particle dimension, and the difference between the second and third dimension, be greater than r_{PSF}.

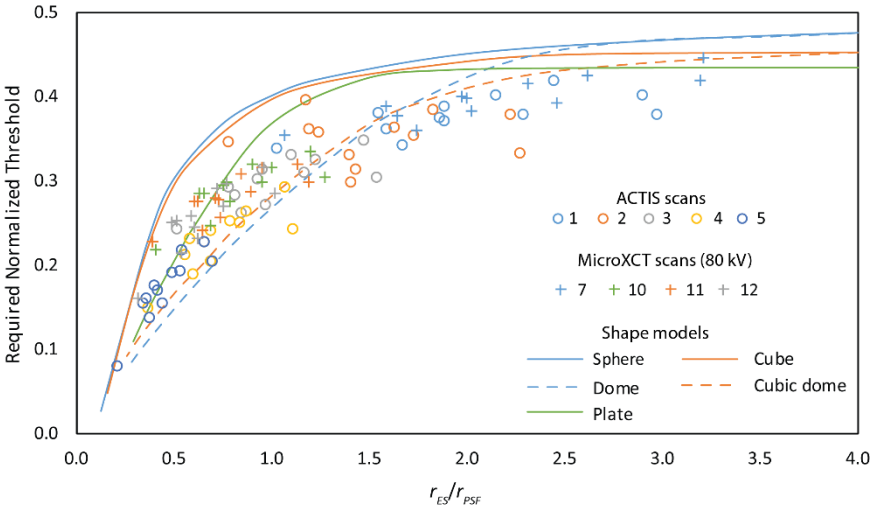

Figure 3. Plot of equivalent-sphere radius (r_{ES}) divided by r_{PSF} against the effective threshold (i.e. the threshold necessary to correctly reproduce the grain volume) normalized to the difference in CT numbers for gold and matrix, for each particle in seven scans. Lines show size-threshold relations predicted by convolving a PSF with idealized shapes. Plate and hollow dome models have thicknesses 1/10 of object width.

5. DISCUSSION

5.1 Robustness and uncertainties

The PVB method is highly resilient against errors in calibration parameters. Unlike thresholds, which can give widely ranging answers when changed only slightly, or omit features entirely if blurring causes them to no longer be within the

threshold range, the uncertainty in PVB-corrected volume is linear with CT_f - CT_{mat} [12]. Thus, an x% error in CT_f - CT_{mat} has only an x% effect on v_{PVB}, which is propagated evenly across all sizes.

Nevertheless, PVB method employs a number of underlying assumptions that should be borne in mind when using it. Chief among these is that attenuation is conserved, and that CT_f is equivalent at all locations in the scan volume. In particular, beam hardening can change CT number meanings both across the scan volume and even within the particles, which, when they become large enough compared to the data resolution, can exhibit bright rims and dark centers. This occurred for the larger grains in this study (Fig. 1, center inset), and is likely to be a factor in the minor resolution dependence observed in PVB results for those grains (Fig. 2). However, the observed error is limited to ~10%, and as gold is among the most challenging natural substances to scan in terms of artifacts, other applications should have less of a problem. However, because of the PVB method's reliance on CT numbers retaining consistent meanings, beam hardening across the sample volume remains a concern.

5.2 First-principles estimation of CT_f

The primary hurdle in employing the PVB method is determining CT_f. Because the X-ray attenuation properties of the materials in this study are relatively well understood, except for the purity of the gold grains, it should be possible to predict the relative attenuation numbers of the grains versus the matrix, opening an avenue to streamlining calibration.

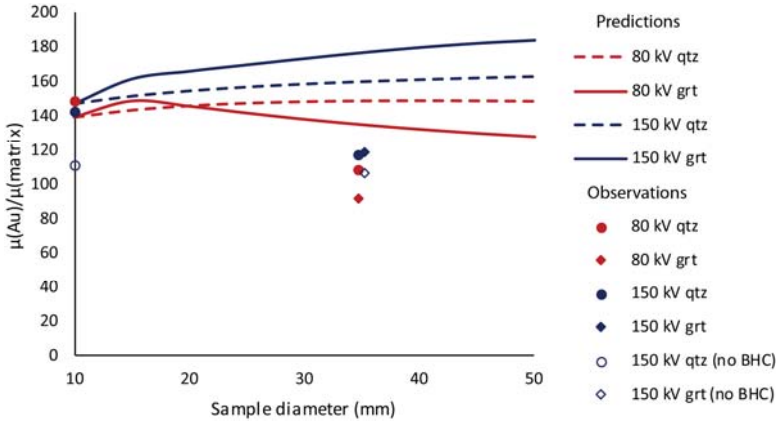

Figure 4. Predicted (lines) versus observed (points) effective attenuation coefficients for gold particles, normalized by the surrounding quartz+epoxy matrix.

The curves in Figure 4 are first-principles estimates of the relative attenuation coefficients of gold versus quartz+epoxy matrix for each of the MicroXCT scans, based on modeled X-ray spectra and their evolution as they pass through the tube chamber, pre-filter, and surrounding garnet and/or quartz, also accounting for the energy dependence of detector efficiency [12]. The prediction is that gold should be ~120-180 times as attenuating as the matrix, with different trends as sample size increases based on the surrounding material, and particularly in how the low- and high-kV spectra interact with the gold K-edge at 80.72 keV.

The plotted points show the results obtained from each scan based on the observed ratio of the CT numbers for gold and matrix, the former obtained via the PVB-method estimate using a known particle. Although the highest-resolution points (10 mm sample diameter) appear to match the predicted values well, recalculation using a revised reconstruction in which the software beam-hardening correction was omitted (open symbol) resulted in sharply lower values. With a 35-μm field of view including quartz or garnet powder, the observed attenuation falls considerably below the prediction, though the effect of eliminating the beam-hardening correction was reduced. Some predicted trends were observed, such as the garnet surrounding medium increasing relative attenuation relative to quartz with a 150 kV spectrum and decreasing it at 80 kV. Nevertheless it is apparent that the theoretical calculation is incomplete [12].

5.3 Practical calibration of CT_f

The previous exercise makes clear that even modest interference with the raw CT data, such as imposing a linearization factor to limit beam-hardening artifacts, can have a large effect on CT numbers, as do sample size and bulk composition. It is thus important that a calibration for CT_f made for a given scan only be applied to very similar scanning arrangements. Acquisition and reconstruction settings must be identical, as should sample geometry and size. Sample composition should

also remain similar. Within these strictures, however, creation and re-use of a calibration can be straightforward. The most practical method is to acquire a scan at the desired conditions, find an example of the phase of interest, and re-scan it at a sufficiently high resolution that resolution effects do not interfere. This may be done using sub-volume scanning, or by cutting a subsample out of the original sample for scanning under more optimal conditions. Once a reliable volume for the feature is obtained, it can be used to calibrate CT_f for all other similar features in the original scan, and all other scans taken at the same conditions.

5.4 Multiple phases

Because the features targeted by the PVB method are by definition small, they can be ambiguous if there are multiple phases that may have similar attenuation properties. Further work on this topic is in progress, but a useful trick is to plot the mean or median CT number for segmented objects against their apparent size after a preliminary threshold. Different phases may cluster in different regions of this plot, allowing them to be discriminated. Each different phase can then be quantified using its respective CT_f value.

6. CONCLUSIONS

A good operating principle for measurement methods is that they should achieve equivalent results across different instruments and resolutions. Altogether, the PVB method for measuring particles close to the resolution limit is accurate, reproducible, resolution-invariant, and objective. It is also notable for its favorable error structure, and is capable of accurately measuring small features below six voxels in volume. In contrast, measuring such features using global thresholds is inherently flawed, as the appropriate threshold changes as a function of feature size and shape, as well as the imaging conditions.

REFERENCES

[1] D. Wildenschild, J. W. Hopmans, C. M. P. Vaz *et al.*, "Using X-ray computed tomography in hydrology: systems, resolutions, and limitations," Journal of Hydrology, 267, 285-297 (2002).

[2] T. M. Ryan, and B. van Rietbergen, "Mechanical significance of femoral head trabecular bone structure in *Loris* and *Galago* evaluated using micromechanical finite element models," American Journal of Physical Anthropology, 126, 82-96 (2005).

[3] ASTM, [E1441-11 Standard Guide for Computed Tomography (CT) Imaging] ASTM International, West Conshohocken, PA(2011).

[4] R. A. Johns, J. S. Steude, L. M. Castanier *et al.*, "Nondestructive measurements of fracture aperture in crystalline rock cores using X-ray computed tomography," Journal of Geophysical Research, 98(B2), 1889-1900 (1993).

[5] A. Keller, "High resolution, non-destructive measurement and characterization of fracture apertures," International Journal of Rock Mechanics and Mining Science, 35(8), 1037-1050 (1998).

[6] M. Van Geet, and R. Swennen, "Quantitative 3D-fracture analysis by means of microfocus X-ray computer tomography (μCT): an example from coal," Geophysical Research Letters, 28(17), 3333-3336 (2001).

[7] S. P. Bertels, D. A. DiCarlo, and M. J. Blunt, "Measurement of aperture distribution, capillary pressure, relative permeability, and in situ saturation in a rock fracture using computed tomography scanning," Water Resources Research, 37(3), 649-662 (2001).

[8] R. A. Ketcham, D. T. Slottke, and J. M. J. Sharp, "Three-dimensional measurement of fractures in heterogeneous materials using high-resolution X-ray CT," Geosphere, 6(5), 499-514 (2010).

[9] S. Prevrhal, K. Engelke, and W. A. Kalender, "Accuracy limits for the determination of cortical width and density: the influence of object size and CT imaging parameters," Physics in Medicine and Biology, 44, 751-764 (1999).

[10] G. M. Treece, A. H. Gee, P. M. Mayhew *et al.*, "High resolution cortical bone thickness measurement from clinical CT data," Medical Image Analysis, 14, 276-290 (2010).

[11] G. M. Treece, K. E. S. Poole, and A. H. Gee, "Imaging the femoral cortex: Thickness, density and mass from clinical CT," Medical Image Analysis, 16, 952-965 (2012).

[12] R. A. Ketcham, and A. S. Mote, "Accurate measurement of small features in X-ray CT data volumes, demonstrated using gold grains," Journal of Geophysical Research - Solid Earth, 124, (2019).

[13] J. R. Kyle, and R. A. Ketcham, "Application of high resolution X-ray computed tomography to mineral deposit origin, evaluation, and processing," Ore Geology Reviews, 65, 821-839 (2015).

[14] F. Reyes, Q. Lin, O. Udoudo *et al.*, "Calibrated X-ray micro-tomography for mineral ore quantification," Minerals Engineering, 110, 122-130 (2017).
[15] R. A. Ketcham, and J. Hildebrandt, "Characterizing, measuring, and utilizing the resolution of CT imagery for improved quantification of fine-scale features," Nuclear Instruments and Methods in Physics Research B, 324, 80-87 (2014).
[16] R. A. Ketcham, "Computational methods for quantitative analysis of three-dimensional features in geological specimens," Geosphere, 1(1), 32-41 (2005).
[17] R. A. Ketcham, and W. D. Carlson, "Acquisition, optimization and interpretation of X-ray computed tomographic imagery: Applications to the geosciences," Computers and Geosciences, 27, 381-400 (2001).
[18] G. T. Herman, "Correction for beam hardening in computed tomography," Physics in Medicine and Biology, 24(1), 81-106 (1979).