Special Section:
Soil Architecture and Function

Ramaprasad Kulkarni
Markus Tuller*
Wolfgang Fink
Dorthe Wildenschild

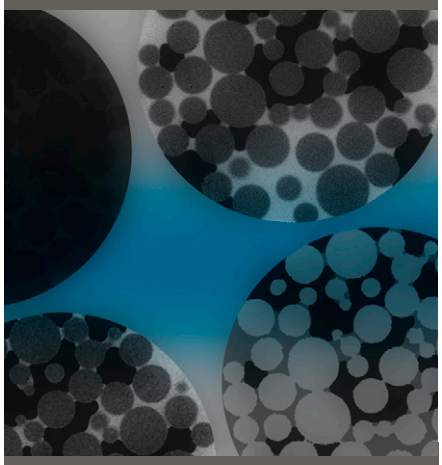

A new method for segmentation of X-ray computed tomography (CT) data of porous materials amenable to true three-dimensional multiphase processing was developed and evaluated. The employed Bayesian Markov random field framework has distinct advantages over commonly applied segmentation algorithms as it can handle any number of phases present in the CT dataset.

R. Kulkarni and W. Fink, Dep. of Electrical and Computer Engineering, Univ. of Arizona, Tucson, AZ 85721; R. Kulkarni and M. Tuller, Dep. of Soil, Water and Environmental Science, Univ. of Arizona, Tucson, AZ 85721; W. Fink, Dep. of Biomedical Engineering, Univ. of Arizona, Tucson, AZ 85721; and D. Wildenschild, School of Chemical, Biological and Environmental Engineering, Oregon State Univ., Corvallis, OR 97331. *Corresponding author (mtuller@cals.arizona.edu).

Vadose Zone J.
doi:10.2136/vzj2011.0082
Received 6 July 2011.

Three-Dimensional Multiphase Segmentation of X-Ray CT Data of Porous Materials Using a Bayesian Markov Random Field Framework

Advancements in noninvasive imaging methods such as X-ray computed tomography (CT) have led to a recent surge of applications in porous media research with objectives ranging from theoretical aspects of pore-scale fluid and interfacial dynamics to practical applications such as enhanced oil recovery and advanced contaminant remediation. While substantial efforts and resources have been devoted to advance CT technology, microscale analysis, and fluid dynamics simulations, the development of efficient and stable three-dimensional multiphase image segmentation methods applicable to large data sets is lacking. To eliminate the need for wet–dry or dual-energy scans, image alignment, and subtraction analysis, commonly applied in X-ray micro-CT, a segmentation method based on a Bayesian Markov random field (MRF) framework amenable to true three-dimensional multiphase processing was developed and evaluated. Furthermore, several heuristic and deterministic combinatorial optimization schemes required to solve the labeling problem of the MRF image model were implemented and tested for computational efficiency and their impact on segmentation results. Test results for three grayscale data sets consisting of dry glass beads, partially saturated glass beads, and partially saturated crushed tuff obtained with synchrotron X-ray micro-CT demonstrate great potential of the MRF image model for three-dimensional multiphase segmentation. While our results are promising and the developed algorithm is stable and computationally more efficient than other commonly applied porous media segmentation models, further potential improvements exist for fully automated operation.

Abbreviations: CT, computed tomography; GUI, graphical user interface; ICM, iterated conditional modes; MAP, maximum a posteriori; MMD, modified Metropolis dynamics; MRF, Markov random field.

In recent years, soil and porous media research has experienced a vast increase in the application of X-ray CT. This trend is attributable to steadily increasing computational capabilities, significant advancements in synchrotron and benchtop micro-CT technology (e.g., Vaz et al., 2011), and easier access to industrial and synchrotron X-ray CT facilities. While substantial efforts and resources have been devoted to advancing CT technology, microscale analysis, and fluid dynamics simulations, the development of adequate three-dimensional, multiphase image segmentation methods for the conversion of inverted grayscale CT volumes into a discrete form that permits quantitative characterization of phase boundaries and subsequent modeling of flow processes and liquid distribution within porous materials seems to lag behind. Although there are numerous segmentation methods documented in the literature (mostly for medical or optical character recognition applications), a stable three-dimensional multiphase algorithm for porous materials is lacking. (For a comprehensive survey and comparison of commonly applied methods, see Iassonov et al. [2009] and Baveye et al. [2010]).

Most of the available codes are limited to two phases (e.g., Lindquist, 1999; Schlüter et al., 2010) or are not applicable to true three-dimensional processing, thereby omitting important information contained in the three-dimensional voxel neighborhood and potentially introducing directional bias (Elliot and Heck, 2007) (note that in contrast to a pixel, which describes the smallest element within a two-dimensional image, a voxel is the smallest element of a three-dimensional data set). Furthermore, many researchers manually define segmentation thresholds based on the shape of the global grayscale or linear attenuation coefficient histograms, which may introduce operator bias in situations where phase contrasts are not clearly defined (i.e., phase attenuation values overlap). It also should be stated that there is a vast discrepancy in segmentation results when different methods are applied

to the same X-ray CT data set (Iassonov et al., 2009; Baveye et al., 2010; Wang et al., 2011).

Due to the lack of multiphase segmentation capabilities, it is common practice to conduct dual-energy or wet–dry scans, segment and align the obtained data sets, and perform subtraction analysis to separate phases (e.g., Porter and Wildenschild, 2010). This is a complex, multistep process susceptible to errors. Very few algorithms developed for medical (Chuang et al., 2006; Held et al., 1997) and pattern and object recognition applications (Berthod et al., 1996) are applicable to direct multiphase segmentation. The fuzzy *c*-means method (Chuang et al., 2006) considers gray-level intensity and probability for each class of voxels (i.e., phase) based on local spatial information calculated for a small voxel neighborhood. Berthod et al. (1996) proposed a two-dimensional algorithm for supervised Bayesian segmentation based on MRFs. Because more nearest neighbors are considered, extension to three dimensions provides more information for the maximum a posteriori (MAP) approach when compared with the two-dimensional model proposed by Berthod et al. (1996). The MAP approach was previously used for image processing, image analysis, and computer vision (e.g., Geman and Geman, 1984; Deng and Clausi, 2005; Sudderth and Freeman, 2008; Chantas et al., 2010; Levada et al., 2010; Paulsen et al., 2010) but has not been applied to three-dimensional multiphase segmentation of porous media X-ray CT data. Although the MAP approach can segment multiple phases, its use has generally been limited to binary segmentation.

Motivated by the documented demand for advanced segmentation capabilities, we developed and implemented a stable and computationally efficient three-dimensional multiphase Bayesian MRF algorithm for segmentation of X-ray CT data of porous materials. Three sample data sets for natural and artificial porous materials emanating from synchrotron microtomography were used to demonstrate the applicability of the algorithm.

Materials and Methods

Code Development and Implementation

Early models for image characterization were based on stochastic representation of image attributes, mainly the gray-level distribution. A widely used model for images based on the probability of classes or gray levels is referred to as the random field model, which was later extended to the MRF model (Moussouris, 1974). The motivation for the application of a stochastic framework is based on the assumption that the variation and interactions among image attributes can be described by probability distributions. The MRF model is inherently powerful for image segmentation because it can generally handle any number of voxel classes (e.g., representing different pore-filling fluids or different solid grain materials); however, it must be initialized with reasonable statistics (i.e., mean and standard deviation) for each voxel class. The statistics can be obtained by either a global thresholding method or simple manual selection of several seed regions within the data that are representative for each voxel class (i.e., phase). The latter method, i.e., manual selection of seed regions, was used in this study. Following the two-dimensional approach of Berthod et al. (1996), we developed and implemented a three-dimensional algorithm optimized for large X-ray CT data sets of porous materials, where each gray level within the three-dimensional data set is assigned a discrete label corresponding to one of the considered individual classes (phases) to satisfy

$$\hat{L} = \arg\min\left\{\sum_{i=1}^{N}\left[\ln\sqrt{2\pi\sigma_L} + \frac{(x_i-\mu_L)^2}{2\sigma_L^2}\right] + \sum_{\{S_i,S_j\}\in C}\beta\gamma\left(l_{S_i},l_{S_j}\right)\right\} \quad [1]$$

with

$$\gamma\left(l_{S_i},l_{S_j}\right) = \begin{cases} -1 & \text{if} \quad l_{S_i} = l_{S_j} \\ +1 & \text{if} \quad l_{S_i} \neq l_{S_j} \end{cases}$$

where l_{S_i} and l_{S_j} are labels for sites S_i and S_j, respectively, corresponding to voxels x_i and x_j in a three-dimensional space, μ_L is the mean and σ_L is the standard deviation of labeling L, N is the total number of voxels within the X-ray CT data set, and β is a constant that represents the homogeneity of regions. A detailed derivation of Eq. [1] is presented in Appendix A. Finding a labeling $\hat{L}$ with the constraint of satisfying Eq. [1] is a combinatorial optimization problem. To investigate different optimization approaches, we initially implemented a deterministic (iterated conditional modes [ICM] algorithm) and a heuristic (Metropolis algorithm) optimization scheme. Because there was only a small adaptation to the Metropolis algorithm required (see Appendix B), we later added and tested the modified Metropolis dynamics (MMD) algorithm.

To facilitate the ease of application, an intuitive graphical user interface (GUI) was developed in MATLAB (The MathWorks, Natick, MA). The GUI provides options for two- or three-dimensional processing, a choice of optimization methods and parameters, tools for data cropping and visualization, and different pre- and post-processing filters. Because MATLAB is based on matrix operations, it lacks efficiency when it comes to the element-by-element computation inherent to the implemented heuristic and deterministic optimization schemes. To assure computational efficiency, all algorithms were coded in the C programming language and interfaced with the GUI via MATLAB executable functions. To further increase efficiency, the Open Multi-Processing (OpenMP) application programming interface that provides compiler directives, functions, and environment variables and enables shared memory parallel programming in C, C++, and Fortran (Chandra et al., 2001) was utilized. An added advantage of OpenMP is its multiplatform support (i.e., UNIX, Linux, Windows, and Mac OS X); C programs with OpenMP directives and functions may be compiled on any of these platforms and the generated executables can utilize multiple cores of a CPU or multiple CPUs in a shared memory system. Locally adaptive segmentation algorithms (Iassonov et al., 2009) such as the presented MRF model are highly suitable for parallel computing, hence utilization of multiple CPUs or CPU cores leads to a significant increase in computation speed. It is also worth

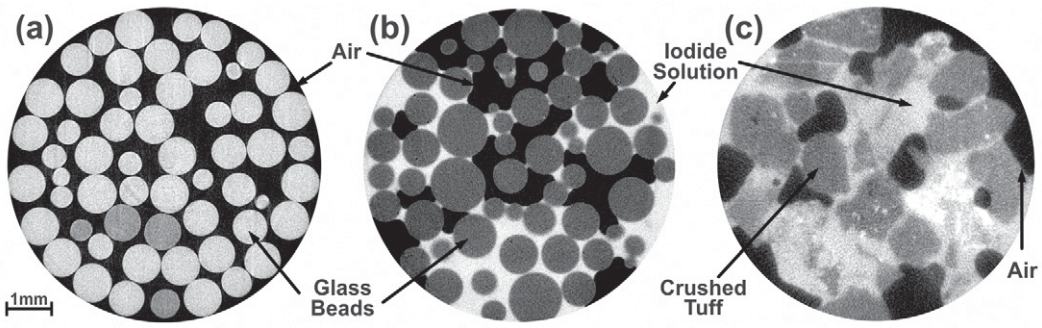

Fig. 1. Samples used for testing the Markov random field code: (a) dry precision glass beads; (b) mixed glass beads and (c) crushed volcanic tuff partially saturated with 11% KI solution (for better visibility of phase boundaries, contrast stretching was applied to original X-ray computed tomography grayscale data).

noting that a parallel random number generator was used for generating random numbers in the parallel implementation of Metropolis and MMD using OpenMP (L'Ecuyer et al., 2002).

To allow input from different data sources, the code was designed to accept various image formats (i.e., jpeg, tiff, bmp, and png) and depth resolutions (i.e., eight-bit, 16-bit, 32-bit, 64-bit, and floating point).

Samples for Model Testing

A total of three X-ray CT data sets were used for MRF code evaluation. The data used emanated from experiments at the synchrotron X-ray microtomography system at the GeoSoilEnviro Consortium for Advanced Radiation Sources (GSECARS) bending magnet beamline at Argonne National Laboratory, Lemont, IL. All details about the GSECARS beamline and scanning procedure were discussed in Porter and Wildenschild (2010) and Wildenschild et al. (2002, 2005).

The first sample (Fig. 1a) was composed of precision glass beads (diameter = 0.8 ± 0.01 mm, density = 2.65 g cm^{-3}, as specified by the manufacturer, Glenn Mills, Clifton, NJ). The second sample (Fig. 1b) consisted of a packing of glass beads of three different sizes (Porter et al., 2010). While the first sample was scanned in a dry state, the second sample was partially saturated (approximately 50%) with an 11% (w/w) KI solution to increase the density contrast between wetting and nonwetting phases. The third sample (Fig. 1c) was composed of crushed volcanic tuff that consisted mainly of quartz with minor amounts of feldspar, albite, and volcanic glasses. This sample was also partially saturated with the 11% KI solution.

All samples were compacted into small acrylic cylinders (i.d. = 6.54 mm) and scanned at an energy level of 33.27 keV at resolutions of 11.8, 13.0, and 16.8 μm. The obtained radiographs were inverted with a filtered backprojection algorithm (Rivers, 2010) and stored in 16-bit grayscale format. The ultimate sizes of the inverted data sets were 650 by 650 by 515 voxels for Sample 1, 650 by 650 by 427 voxels for Sample 2, and 650 by 650 by 380 voxels for Sample 3. For the dry glass bead sample (Fig. 1a), which was primarily used as a benchmark for model testing, the porosity (ϕ_M = 0.509) was independently determined based on the number, specific gravity, and mass of precision glass beads required to fill the cylindrical sample container. The container dimensions for volume calculation were precisely determined from a vertical radiograph cross-section, i.e., to determine sample height, the container was truncated at the lower and upper edges of the lowest and uppermost glass bead, respectively.

Data Preprocessing

To omit container walls and the surrounding air space that were captured together with the actual porous medium sample, all grayscale X-ray CT data sets needed to be cropped before segmentation using a cylindrical mask (Fig. 2a). To assure that identical data sets were used for sensitivity testing of model input parameters, the cropped volumes were saved and called for all subsequent enhancement and segmentation steps.

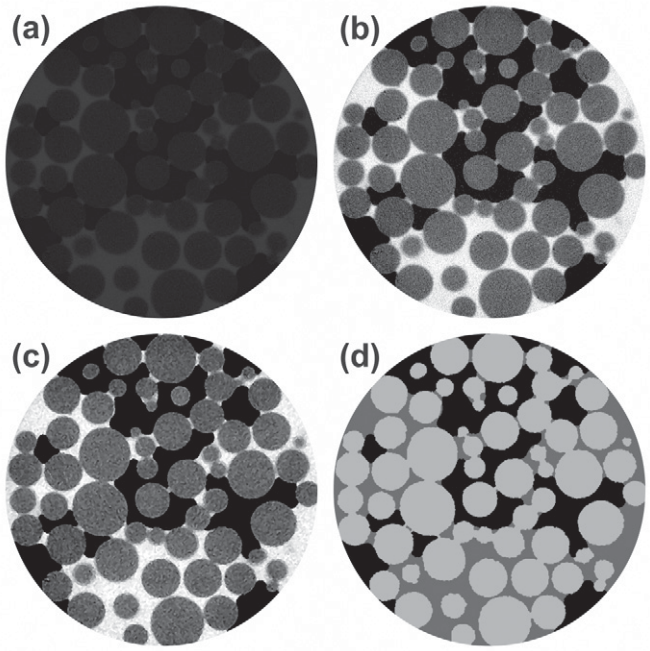

Fig. 2. (a) Cropped example cross-section of the original grayscale data set of the partially saturated glass bead sample, the same cross-section after (b) contrast stretching and (c) application of median and unsharp filters, and (d) segmented cross-section.

As shown for an example cross-section of the partially saturated glass bead sample (Fig. 2a), grayscale X-ray CT data emanating from synchrotron sources appear very dark on a computer screen, with little contrast among phases. This means that the global grayscale distribution is rather narrow, occupying only a small portion of the 16-bit (0–65,535 grayscale levels) range. This can be significantly improved without the loss of physical information via contrast stretching, where the narrow distribution is linearly mapped (stretched) across the entire 16-bit grayscale range. This yields brighter images with better phase contrast (Fig. 2b) and allows a better informed manual selection of seed regions.

To remove high-frequency noise inherent to X-ray CT data and to enhance phase boundaries, we applied median and unsharp filters (Joshi, 2006) to the contrast-stretched data sets (Fig. 2c). The median filter was applied in a 3- by 3- by 3-voxel window. The choice of a small voxel neighborhood is imperative for preserving phase edges while removing noise. The unsharp filter enhances phase boundaries locally by subtracting a smoothed or unsharp version of an image from its original to obtain an edge or high-frequency image, which is subsequently added back to the original image.

After image enhancement, seed regions for each phase were selected manually from random locations within the three-dimensional data set. To accomplish this task, we randomly selected at least three two-dimensional cross-sections close to the bottom, center, and top of the sample and marked seed regions within each considered phase until the volume fraction of seed voxels was at least 7.50×10^{-4}. Because of the monochromatic nature of the synchrotron X-ray source, which yields uniform phase densities (i.e., no beam hardening), seeding from a single cross-section would theoretically be sufficient. For data sets from polychromatic sources (e.g., industrial or benchtop X-ray CT scanners), where beam hardening is an issue, an iterative procedure (Iassonov and Tuller, 2010) can be applied to remove density variations prior to segmentation.

To assure that identical seed regions were used for sensitivity testing of the input parameters and comparison of the implemented optimization schemes, the mean and standard deviation of the grayscale distribution of each selected seed region were calculated and stored in a separate file and later called for the subsequent segmentation step. After seed region selection, the MRF and optimization parameters were specified on the GUI and the segmentation step was initialized. After segmentation, the resulting data set is stored in TIFF format and is ready for further analysis (e.g., determination of porosity). An example cross-section for the segmented partially saturated glass bead data set is depicted in Fig. 2d.

Results and Discussion

All segmentation results presented here were obtained with a Windows 7 (64-bit) workstation with an Intel Core i7–980X Extreme (12-MB L2 cache, 3.33 GHz) CPU with six physical cores (12 threads with hyper-threading), 24-GB DDR3 SDRAM at 1333 MHz, and serial ATA II RAID 1 dual 1-TB hard drives.

Sensitivity of Segmentation Results to Statistical Initialization

Unlike for biomedical applications (Held et al., 1997), where the statistics of all image classes of interest are known a priori, the MRF image model for porous materials needs to be initialized with the means and standard deviations of the grayscale distributions of all considered classes (phases), as discussed above and presented in detail in Appendix A. In this study, we applied manual seeding by randomly selecting regions within each phase as outlined above. To evaluate the effects of the initial seeding on the segmentation results, we gradually increased the volume fraction of seed voxels for each phase for the partially saturated glass bead and crushed tuff samples (three-phase systems) to approximate the minimum fraction required for stable and repeatable segmentation results. Figure 3a shows that for glass beads, the segmentation results stabilized when the volume fraction used for the statistical initialization of each phase was at least 6.07×10^{-4}. For crushed tuff (Fig. 3b), the cutoff volume fraction (marked with a dashed line) was around 7.50×10^{-4}. Note that the solid phase (i.e., glass beads and tuff) was assigned a separate *y* axis to visualize the variations in greater detail (Fig. 3).

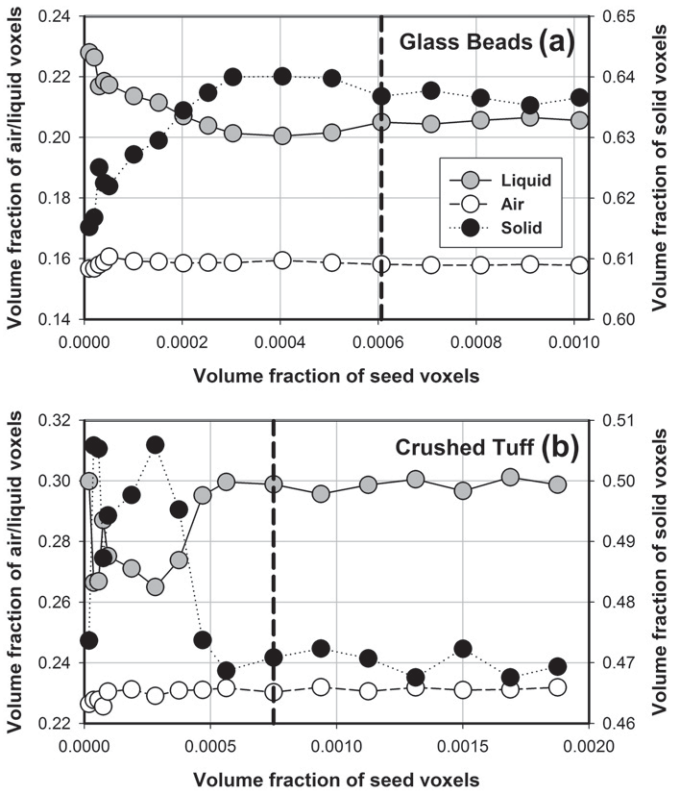

Fig. 3. Volume fractions of voxels for (a) partially saturated glass beads and (b) crushed tuff plotted against the volume fraction of seed voxels. Results were obtained with iterated conditional modes optimization with Markov random field β set to 0.9.

Although there are small variations visible beyond the cutoff thresholds, the obtained results are very solid. If the sample porosity is expressed as ϕ_{CT} = (no. of air voxels + no. of liquid voxels)/(total no. of voxels), the variance for porosities obtained for the glass beads for the seed voxel volume fraction range between 6.07×10^{-4} and 1.01×10^{-3} was only 7.6×10^{-7} (Fig. 3a). Considering the seed voxel volume fraction range between 7.50×10^{-4} and 1.88×10^{-3} for the crushed tuff (Fig. 3b), the variance for the calculated porosities was 4.0×10^{-6}. Based on these results, a minimum seed region volume fraction of 7.50×10^{-4} was chosen for all subsequent segmentations in this study.

It should be noted that manual seeding is very time consuming, with time requirements often exceeding the actual computation time. While manual seeding provides reasonably stable and repeatable segmentation results, we acknowledge the need for fully automated and unbiased statistical initialization. This could be achieved with a two-step approach, where, for example, a simple and computationally efficient clustering method (Iassonov et al., 2009) is utilized to automatically detect and classify the phases of interest before application of the MRF model. The implementation of such an automation scheme is rather complex and part of our ongoing research.

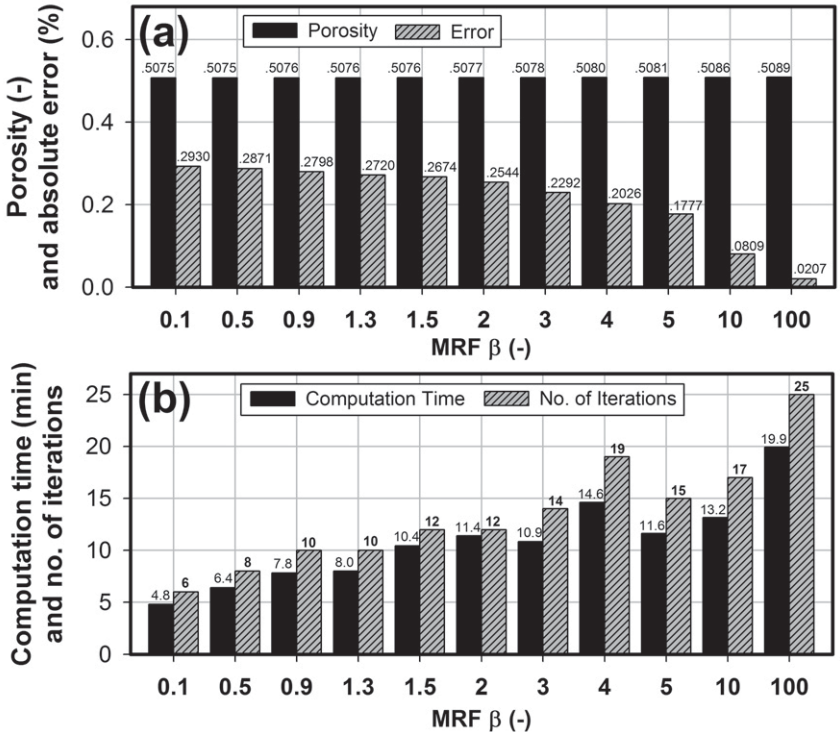

Fig. 4. (a) X-ray computed tomography derived porosities and associated absolute percentage errors for the dry glass bead sample computed for a range of values of the Markov random field (MRF) β, and (b) associated computation times and number of iterations required for model convergence. Results were obtained with iterated conditional modes optimization.

Sensitivity of Segmentation Results to Markov Random Field Model Parameterization

Besides initialization statistics for all considered phases, the MRF image model requires parameterization with a β value (Eq. [1]) that emphasizes the homogeneity of regions within the X-ray CT data set. The choice of β reflects a priori assumptions about the expected variation among the respective phases. A small β value is chosen for spatially more heterogeneous phase distributions, whereas a large β is chosen for more homogeneous conditions (note that β is a positive number). The choice of β not only affects the segmentation results but also determines the time it takes for the algorithm to converge (i.e., $\Delta E < E_{threshold}$; see Appendix B). To evaluate the effects of MRF model parameterization on the segmentation results and computational efficiency, we first gradually increased β from 0.1 to 100 for the dry glass bead sample and compared the X-ray CT derived porosities with the physically measured value of $\phi_M = 0.509$ (see above). Figure 4a depicts the X-ray CT derived porosities and associated absolute percentage error $E_A = |[(\phi_{CT} - \phi_M)/\phi_M]100|$ plotted against β. For all considered β values, the X-ray CT derived porosities were very close to the benchmark value, with E_A gradually decreasing with increasing β. Considering the noise in the original X-ray CT data set, even the largest deviation, with $E_A = 0.293\%$ for β, seems to be reasonable for subsequent morphometric pore space analysis and fluid dynamics modeling. While the CT-derived porosities showed only small variations, the associated computation times and the number of iterations required for model convergence varied significantly (Fig. 4b). Although there was no gradual trend, the computation times and number of iterations were considerably larger for higher β values.

The sensitivity of the segmentation results to changes in β was also qualitatively evaluated for the partially saturated glass bead and crushed tuff samples (Fig. 5). For these data sets, we unfortunately did not have a benchmark from independent measurements for comparison. Nevertheless, the obtained results for the partially saturated glass beads (Fig. 5a) showed a very similar trend as for the dry glass beads (Fig. 4). There was virtually no change in porosity across the entire range of considered β values. The numbers of segmented phase voxels are depicted in Fig. 5a. The variance of the CT-derived porosities, ϕ_{CT}, across the observed β range was only 1.5×10^{-6}. The crushed tuff sample showed vastly different behavior (Fig. 5b). While the CT-derived porosities were virtually constant up to β = 1.5, starting at β = 2.0 the MRF model misclassified the solid and liquid phases, with the air phase remaining virtually constant. This can be attributed to the poorly pronounced solid–liquid interfaces in the original crushed tuff CT data set (Fig. 1c), which could be potentially improved with more sophisticated

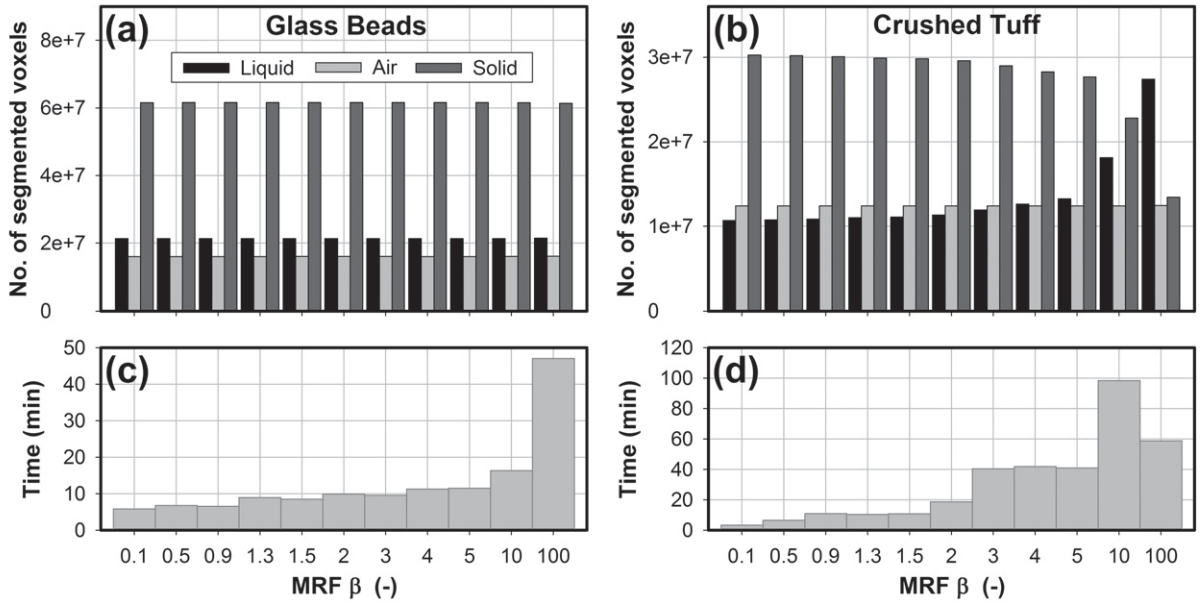

Fig. 5. Number of segmented liquid, air, and solid voxels for (a) partially saturated glass beads and (b) crushed tuff for a range of values of the Markov random field (MRF) β, and associated computation times for (c) glass beads and (d) tuff. Results were obtained with iterated conditional modes optimization.

image enhancement techniques such as the anisotropic–nonlinear diffusion filter (Kaestner et al., 2008). A visual comparison of the segmentation results obtained for the crushed tuff with different β values is shown in Fig. 6. As observed for the dry glass beads, the computational demand was significantly higher for large β values for both the partially saturated glass beads and the tuff (Fig. 5c and 5d). Based on the above analysis and considering the tradeoff between computational efficiency and accuracy, β values between 0.5 and 1.3 are reasonable for segmenting X-ray CT data such as used for this study. Hence, a value of 0.9 (arithmetic mean) was chosen for all subsequent segmentations presented below.

Comparison of Implemented Combinatorial Optimization Schemes

Before proceeding with a direct comparison of the optimization methods used (i.e., Metropolis, MMD, and ICM optimization), the required parameterization of the Metropolis and MMD algorithms with initial temperature T_0, cooling schedule c, and relabeling constant α (only used for MMD; see Appendix B) need to be discussed. While the cooling schedule has been intensively studied, with theoretical limits and formulations proposed in the literature (Kirkpatrick et al., 1983; Yip and Pao, 1995, and references therein), there is no universal criterion for the choice of T_0. The cooling schedule ($0 < c < 1$) determines the time required for the algorithm to converge. Based on published data, c was set at 0.98, which means a 2% temperature reduction after each iteration step. While the choice of a high T_0 may lead to a significant increase in computation time, the selection of a low T_0 might not provide enough time and freedom for the relabeling operation to approach the global minimum energy, instead yielding a local minimum. To evaluate the effects of T_0 on the segmentation results and computational efficiency, we gradually increased T_0 from 0.1 to 20 for the dry glass bead sample and compared the ϕ_{CT} values

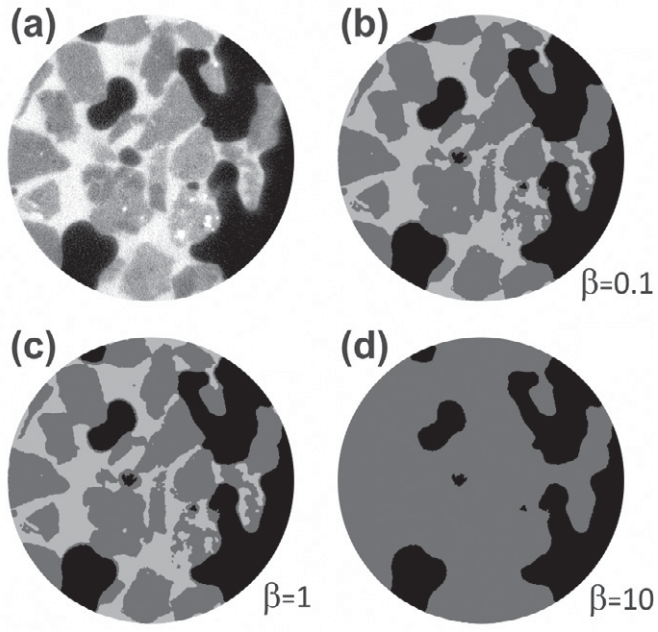

Fig. 6. Visual comparison of segmentation results for crushed tuff obtained with different values of the Markov random field β: (a) original contrast-stretched sample cross-section, and segmented sample cross-sections obtained with (b) β = 0.1, (c) β = 1, and (d) β = 10. Results were obtained with iterated conditional modes optimization.

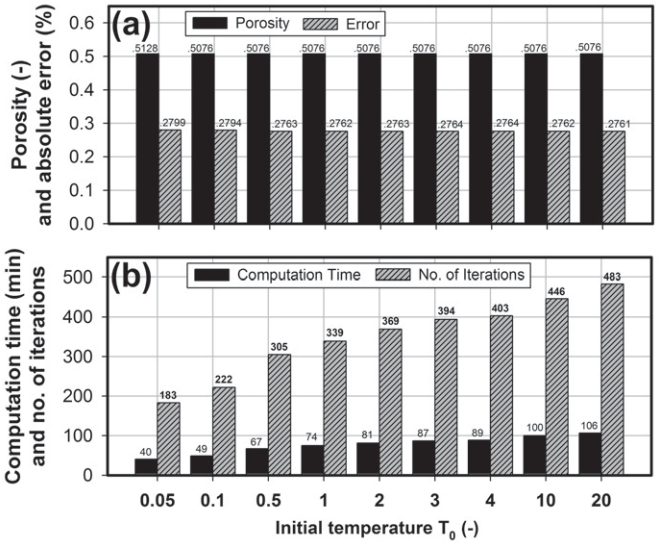

Fig. 7. (a) X-ray computed tomography derived porosities and associated absolute percentage errors for the dry glass bead sample computed for a range of initial temperature (T_0) values, and (b) associated computation times and number of iterations required for model convergence. Results were obtained with Metropolis optimization (cooling schedule c = 0.98) and Markov random field β set to 0.9.

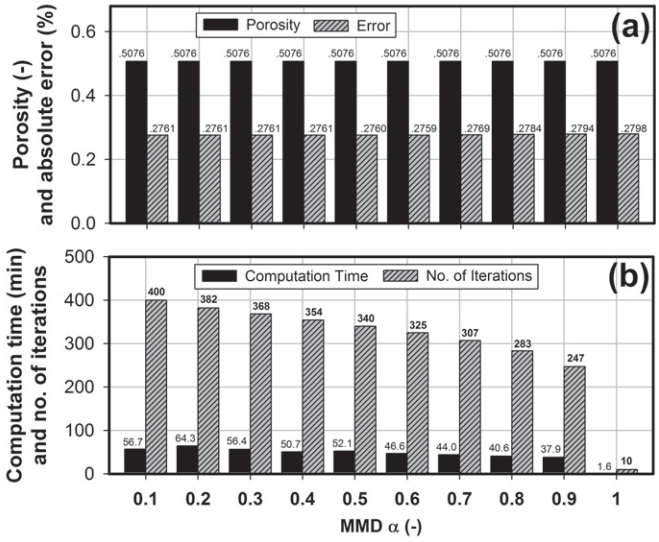

Fig. 8. (a) X-ray computed tomography derived porosities and associated absolute percentage errors for the dry glass bead sample computed for a range of relabeling constant α values for the modified Metropolis dynamics (MMD) algorithm, and (b) associated computation times and number of iterations required for model convergence. Results were obtained with initial temperature T_0 = 1, cooling schedule c = 0.98, and Markov random field β set to 0.9.

with the physically measured porosity, ϕ_M (Fig. 7). Because of the similarity of the Metropolis and MMD optimization schemes, only the Metropolis algorithm with a cooling schedule of c = 0.98 was evaluated. As evident from Fig. 7a, the choice of T_0 had no impact on the segmentation results for the dry glass beads, which slightly underestimated ϕ_M with an absolute percentage error of about 0.28%; however, there was a significant gradual increase in the number of required iterations and associated computation times with increasing T_0. Based on the forgoing evaluation, we chose T_0 = 1 for all subsequent segmentations with Metropolis and MMD optimization in this study.

As detailed in Appendix B, the only difference between Metropolis and MMD optimization lies in the condition for accepting a new labeling. While this is done randomly in the Metropolis algorithm, MMD uses a deterministic condition. Therefore, an additional relabeling constant α needs to be specified if the MMD optimization is applied. As with T_0, the choice of α had no significant impact on the segmentation results for the dry glass beads. The absolute percentage error for all considered α values was about 0.28% (Fig. 8a). In contrast to T_0, the number of required iterations decreased with increasing α (Fig. 8b), which can be expected based on the MMD pseudocode shown in Appendix B. For the comparison of optimization schemes, α = 0.5 was chosen for the MMD algorithm.

As shown in Fig. 9a for the dry glass beads, the results obtained with Metropolis, MMD, and ICM optimization are indistinguishable. The CT derived porosities of ϕ_{CT} = 0.5076 are identical to the fourth digit, with the ICM algorithm being by far the most computationally efficient (Fig. 9b). Similar results were obtained for the partially saturated glass beads and crushed tuff (Fig. 10a and 10b). The CT-derived porosities for Metropolis, MMD, and ICM were identical to the third digit (ϕ_{CT} = 0.378) for the glass beads. The ϕ_{CT} values for the crushed tuff were 0.431 for Metropolis and MMD and 0.436 for ICM. Again, the ICM algorithm was superior with regard to computational efficiency. Based on this analysis, it is safe to select ICM as the optimization scheme of choice for solving the MRF labeling problem as stated in Eq. [1].

Computational Efficiency of the Implemented Markov Random Field Framework

To discuss the computational efficiency of the implemented MRF algorithm, we consider ICM optimization only because it is the most efficient and most likely scheme to be used for solving the optimization problem stated in Eq. [1]. It also needs to be noted upfront that the presented processor times are inherent to the workstation that was used for this study; nonetheless, they provide some general information about computational efficiency. It is obvious that the segmentation of large X-ray CT data sets is computationally demanding, with the number of computations increasing in an approximately linear fashion with increasing sample size (i.e., number of voxels), as shown for the partially saturated glass beads and tuff in Fig. 11. Note that the number of voxels was gradually increased by expanding the sample height until the full height was reached. The processing times for the entire glass bead (650 by 650 by 427 voxels) and crushed tuff (650 by 650 by 380 voxels) samples were 6.53 and 10.83 min, respectively.

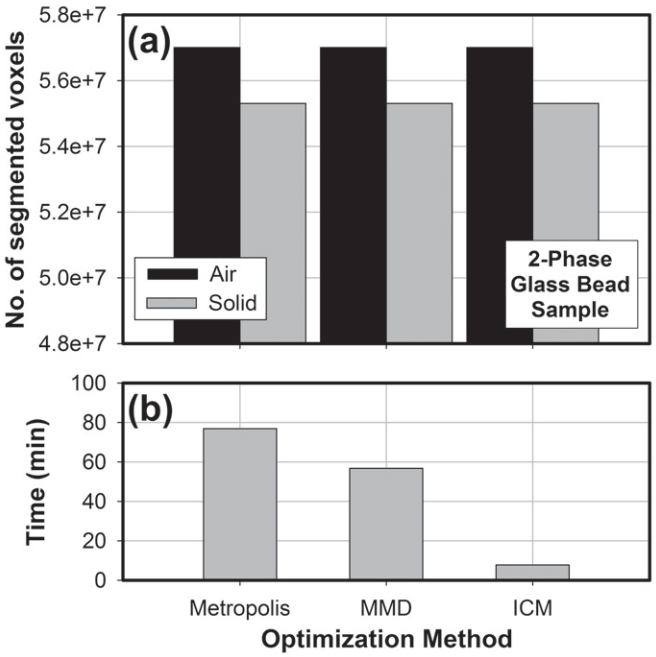

Fig. 9. (a) Comparison of segmentation results for dry glass beads obtained with Metropolis, modified Metropolis dynamics (MMD), and iterated conditional modes (ICM) optimization, and (b) associated computation times. Results were obtained with initial temperature $T_0 = 1$, cooling schedule $c = 0.98$, MMD relabeling constant $\alpha = 0.5$, and Markov random field β set to 0.9.

While the results depicted in Fig. 11 imply a linear increase in computation time with increasing sample size, for very large X-ray CT data sets linearity might be lost because the overhead due to the creation and release of slave threads increases with the number of iterations required for model convergence. For a 2000- by 2000- by 951-voxel, three-phase data set (which we could not utilize for this study due to a confidentiality clause), the computation time was only 58 min. Based on the comprehensive experience of researchers with other global and locally adaptive segmentation algorithms that were tested with similar hardware (Iassonov et al., 2009), the implemented MRF algorithm is reasonably fast, especially when compared with other locally adaptive methods.

Another avenue to increase computational speed, which is part of our ongoing research, is the utilization of graphics processing units (GPUs) that traditionally were confined to video cards and have been mostly used to accelerate the processing and building of images. Recent technological advances have led to the development of general purpose graphics processing units (GPGPUs), which are perfectly suited for parallel computing due to their large number of physical cores that can perform particular floating-point operations, unlike a CPU, which performs general purpose operations. The GPGPUs cannot be used as standalone units because they require initial instructions and data sent from a CPU to carry out parallel operations on a few hundred cores and then return the results to the CPU. Preliminary tests with a NVIDIA Tesla GPU computing system (e.g., D870 with two GPUs, 3-Gb dedicated memory, 860 GFLOPS) using the CUDA SDK (NVIDIA Corp.,

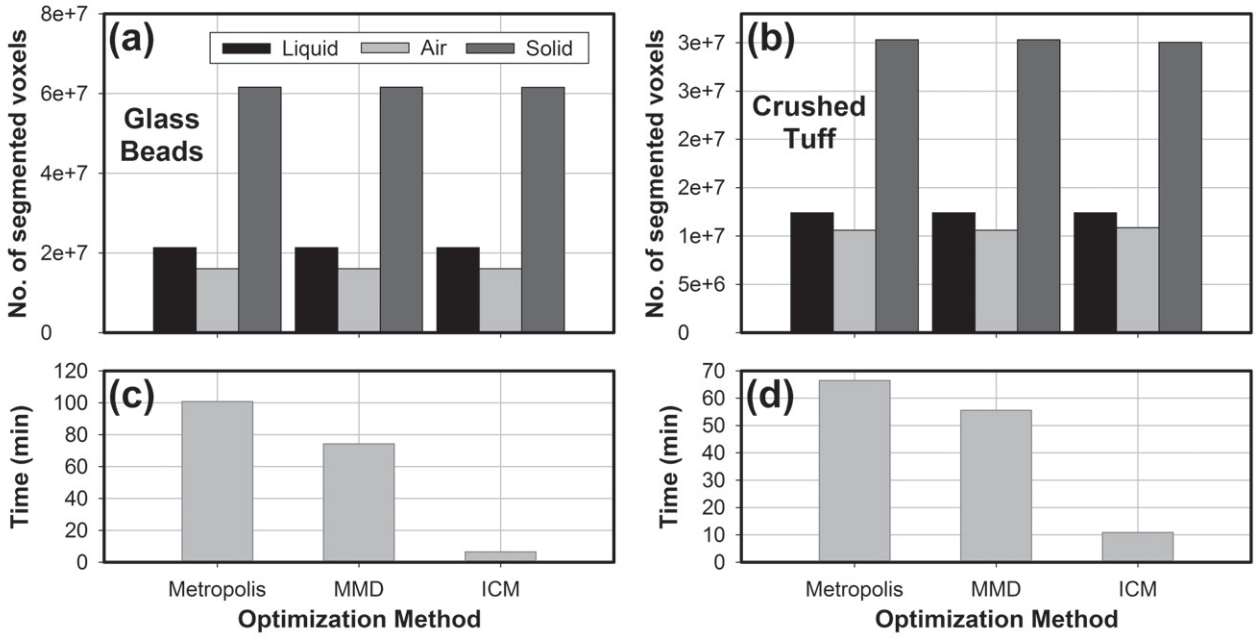

Fig. 10. Comparison of segmentation results for partially saturated (a) glass beads and (b) crushed tuff obtained with Metropolis, modified Metropolis dynamics (MMD), and iterated conditional modes (ICM) optimization and associated computation times for (c) glass beads and (d) tuff. Results were obtained with initial temperature $T_0 = 1$, cooling schedule $c = 0.98$, MMD relabeling constant $\alpha = 0.5$, and Markov random field β set to 0.9.

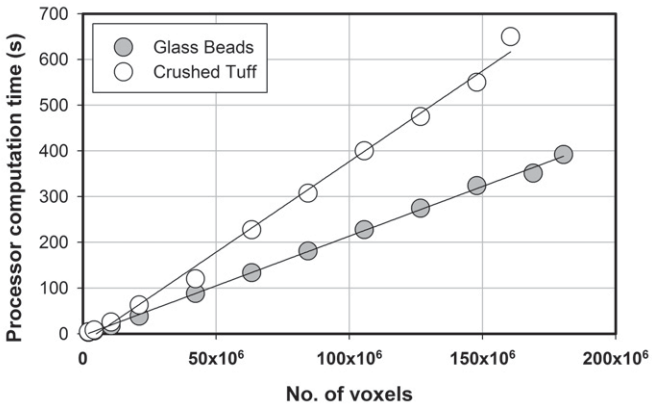

Fig. 11. Processor computation times plotted as a function of the number of image voxels for partially saturated glass beads and crushed tuff. Results were obtained using iterated conditional modes optimization with Markov random field β set to 0.9.

2007) programming language yielded promising results with the potential of a 100-fold increase in computation speed for the samples presented in this study.

Summary and Conclusions

A multiphase segmentation method based on a Bayesian MRF framework amenable to true three-dimensional processing was developed and evaluated. Two heuristic (Metropolis and MMD) and one deterministic (ICM) combinatorial optimization schemes required to solve the labeling problem of the MRF image model were implemented and tested for computational efficiency and their impacts on the segmentation results. The test results for three grayscale data sets consisting of dry glass beads, partially saturated glass beads, and partially saturated crushed tuff obtained with synchrotron X-ray micro-CT demonstrate the great potential of the MRF image model for three-dimensional multiphase segmentation. Evaluation of the manual statistical seeding used in this study revealed that for the investigated X-ray CT data sets, selection of at least 6×10^4 voxels from each considered phase was required for stable and repeatable results. Evaluation of the effects of the MRF β parameter, which defines the homogeneity of phases within the X-ray CT data set, on the segmentation results and computational efficiency yielded an applicable range bound by 0.5 and 1.3. Values of $T_0 = 1$ for the initial temperature and $c = 0.98$ for the cooling schedule, parameters inherent to Metropolis and MMD optimization, were derived considering relative percentage errors and computation time. A comparison of all implemented optimization schemes showed that the ICM algorithm was by far the most computationally efficient and yielded almost identical results to Metropolis and MMD.

Besides computational efficiency and stability, the most distinct advantage of the MRF image model is that it can handle any number of voxel classes (e.g., representing different pore-filling fluids or different solid grain materials), eliminating the need for wet–dry or dual-energy scans, image alignment, and subtraction analysis commonly applied in X-ray micro-CT analysis. Furthermore, the MRF model was implemented for full three-dimensional processing, which is a significant advantage when compared with many other codes that only allow two-dimensional "slice by slice" processing of X-ray CT data.

To eliminate potential operator bias, we are currently working on a fully automated version that utilizes a simple and efficient segmentation method for automated seed region selection and initialization of the MRF model. For further increase in computation speed, the current code is converted to CUDA to take advantage of the powerful parallel computing environment of GPGPUs.

Appendix A

Three-Dimensional Bayesian Markov Random Field Model for Multiphase Segmentation

The applicability of MRF theory to image modeling is based on the Hammersley–Clifford theorem (Besag, 1974) that links MRF and Gibbs probability distributions. Because of the complexity of deriving a joint probability distribution P_X for MRF from its conditional distribution, the theorem is applied to reduce P_X to the simpler Gibbs distribution (Kindermann and Snell, 1980; Li, 2009). Furthermore, the local property of MRF yields algorithms that can be implemented in a massively parallel manner, making MRF a reliable method for solving the image segmentation problem.

The theory of MRF on graphs was first proposed by Geman and Geman (1984). Bayesian MRF image segmentation utilizes graph theory (Berthod et al., 1996) by assigning a discrete set of sites S, which are equivalent to the voxel locations within the X-ray CT data set, and a discrete set of labels that correspond to the individual phases or classes, $L = (1, 2, ..., m)$, within the X-ray CT data set. Considering n sites and m labels yields m^n possible labeling options.

Representing the gray-level attribute of an image voxel with x_i defines an image as $X = (x_i, 1 \leq i \leq N)$, where N is the total number of voxels in the X-ray CT data set. Applying the degraded image model (Geman and Geman, 1984; Derin et al., 1984), the segmentation task is to find a label L for the voxel set X that maximizes the conditional probability $P(L|X)$. Applying Bayes' theorem, the conditional probability can be expressed as

$$P(L|X) = P(X|L)\frac{P(L)}{P(X)} \qquad [A1]$$

Because additive white Gaussian noise is assumed (i.e., the random variables, representing voxel values $[x_i, 1 \leq i \leq N]$ within the image X are independent and identically distributed), $P(X|L)$ may be written as

$$P(X|L)=\prod_{i=1}^{N}P(x_i|L) \qquad [A2]$$

where $P(x_i|L)$ is given by the Gaussian distribution

$$P(x_i|L)=\frac{1}{\sqrt{2\pi\sigma_L}}\exp\left(\frac{-(x_i-\mu_L)^2}{2\sigma_L^2}\right) \qquad [A3]$$

with μ_L as the mean and σ_L as the standard deviation of labeling L. The conditional probability $P(X|L)$ is the probability of observing the respective grayscale values for a voxel set X given a particular labeling L of the entire three-dimensional volume, and $P(x_i|L)$ is the probability of observing the grayscale value of the individual voxel x_i given the labeling L of the entire three-dimensional volume.

Following the two-dimensional framework of Berthod et al. (1996), the Markovian term in Eq. [A1] that describes the influence of neighboring voxels was extended to three dimensions:

$$P(L)=\exp\left[-\frac{1}{T}\sum_{\{S_i,S_j\}\in C}\beta\gamma\left(l_{S_i},l_{S_j}\right)\right] \qquad [A4]$$

where

$$\gamma\left(l_{S_i},l_{S_j}\right)=\begin{cases}-1 & \text{if} \quad l_{S_i}=l_{S_j}\\ +1 & \text{if} \quad l_{S_i}\neq l_{S_j}\end{cases}$$

with l_{Si} and l_{Sj} as labels for sites S_i and S_j, respectively, corresponding to voxels x_i and x_j in a three-dimensional space, and β as a constant that represents the homogeneity of individual phases.

Because the joint probability distribution $P(X)$ for observing the grayscale values for a voxel set X is a constant $0 \leq \varepsilon \leq 1$ independent of labeling, the a posteriori probability $P(L|X)$ of Eq. [A1] is maximized following Eq. [A2]:

$$P(L|X)\propto\prod_{i=1}^{N}P(x_i|L)P(L) \qquad [A5]$$

Considering Eq. [A3] and [A4], Eq. [A5] reduces to

$$P(L|X)=\prod_{i=1}^{N}\frac{1}{\sqrt{2\pi\sigma_L}}\exp\left[\frac{-(x_i-\mu_L)^2}{2\sigma_L{}^2}\right]\times\exp\left[-\frac{1}{T}\sum_{\{S_i,S_j\}\in C}\beta\gamma\left(l_{S_i},l_{S_j}\right)\right] \qquad [A6]$$

Following the MAP labeling solution presented by Berthod et al. (1996) yields

$$\hat{L}=\arg\max\left\{\frac{1}{T}\ln\left[P(X|L)\right]+\ln\left[P(L)\right]\right\}$$

$$\hat{L}=\arg\max\left\{\sum_{i=1}^{N}-\frac{1}{T}\left[\ln\sqrt{2\pi\sigma_L}+\frac{(x_i-\mu_L)^2}{2\sigma_L^2}\right]+\sum_{\{S_i,S_j\}\in C}-\frac{1}{T}\beta\gamma\left(l_{S_i},l_{S_j}\right)\right\} \qquad [A7]$$

With $-1/T < 0$, Eq. [A7] reduces to

$$\hat{L}=\arg\min\left\{\sum_{i=1}^{N}\left[\ln\sqrt{2\pi\sigma_L}+\frac{(x_i-\mu_L)^2}{2\sigma_L^2}\right]+\sum_{\{S_i,S_j\}\in C}\beta\gamma\left(l_{S_i},l_{S_j}\right)\right\} \qquad [A8]$$

Note that there are two terms in Eq. [A8]. The first term is due to the value at the site or gray level for a particular labeling L, and the second term is due to the interactions with neighboring voxels. Bayesian MRF segmentation is basically the assignment of labels to image voxels such that Eq. [A8] is satisfied. Finding a labeling with the constraint of satisfying Eq. [A8] is a combinatorial optimization problem. Many heuristic optimization schemes (for an overview, see Holland, 1975; Goldberg, 1989; Glover and Kochenberger, 2003; Blum and Roli, 2003; and references therein) and some deterministic optimization schemes (e.g., Geman and Geman, 1984; Besag, 1986) are documented in the literature. As discussed above, we strove to test and evaluate heuristic as well as deterministic optimization schemes, as detailed in Appendix B.

Appendix B
Implemented Heuristic and Deterministic Optimization Schemes

The labeling problem described in Appendix A may be solved with combinatorial optimization. To investigate fundamentally different approaches, we initially implemented a deterministic (ICM algorithm) and a heuristic (Metropolis algorithm) optimization scheme. Because there was only a small adaptation to the Metropolis algorithm required, we later added and tested the MMD algorithm, as detailed below.

Metropolis Optimization Algorithm

Metropolis optimization (Metropolis et al., 1953; Kirkpatrick et al., 1983) is a modified Monte Carlo method used to obtain a sequence of random samples from a probability distribution when

direct sampling is challenging. The Metropolis algorithm is suitable for cases where the objective is to obtain an acceptable solution within a reasonable time while avoiding exhaustive searches for the best solution without any time constraint.

The principle behind this algorithm is to calculate the properties of any substance that can be considered as a composition of interacting individual molecules. For the case of image segmentation, the objective is to find an acceptable labeling that satisfies Eq. [A8]. This requires defining an "energy" function (analogous to the interaction energy of the molecules within a substance) that represents the current labeling, which needs to be minimized by randomly choosing new labels. From Eq. [A8], the argument of the argmin( ) term becomes that energy function. The pseudocode of the implemented Metropolis algorithm is

1. Initialize a labeling $\hat{L}_{current} \leftarrow \hat{L}_0$ for which the energy $E_{current} \leftarrow E(\hat{L})$ due to the first term in Eq. [A8] is a minimum, and initialize temperature $T_{current} \leftarrow T_0$.
2. Repeat:
   i. randomly assign a new label: $\hat{L}_{new} \leftarrow$ neighbor ($\hat{L}_{current}$)
   ii. calculate the energy due to the new label: $E_{new} \leftarrow E(\hat{L}_{current})$
   iii. if $\exp[-(E_{new} - E_{current})/T_{current}] >$ random ( ), then accept label and energy $\hat{L}_{new} \leftarrow \hat{L}_{current}$, $E_{new} \leftarrow E_{current}$, respectively. If not, go to 2.i.
   iv. calculate the change in energy: $\Delta E \leftarrow E_{new} - E_{current}$
   v. calculate a new temperature and update: $T_{new} \leftarrow T_{current} \times c$

   until $\Delta E < E_{threshold}$

In Step 1, a labeling $\hat{L}_0$ for all voxels, which corresponds to the least energy contribution due to all possible labels for each voxel, l_i, is assigned and its corresponding energy is calculated. Also an initial temperature is chosen. In Steps 2.i and 2.ii, the voxels are randomly assigned new labels individually, resulting in a new labeling of the entire three-dimensional volume, $\hat{L}_{new}$, and a corresponding energy, E_{new}. In Step 2.iii, the new labeling is accepted if $\exp[-(E_{new} - E_{current})/T_{current}$ is greater than a random number (in the range [0,1])—the Boltzmann probability condition for the Metropolis algorithm. In Step 2.iv, the change in energy is determined. The loop continues after reducing the temperature ("cooling"). These steps are repeated until the change in energy is below a predefined threshold, i.e., the exit criterion is reached.

Because of the similarity to the Metropolis algorithm, we also implemented the MMD algorithm (Kato et al., 1992). The only difference lies in the condition for accepting a new labeling. While this is done randomly in the Metropolis algorithm, MMD uses a deterministic condition. Therefore, only the *if* statement in the above pseudocode needs to be replaced with: if $\exp[-(E_{new} - E_{current})/T] >$ const $(0 \leq \alpha \leq 1)$.

Iterated Conditional Modes Algorithm

Unlike the Metropolis and MMD algorithms, the ICM algorithm is a deterministic optimization scheme (Besag, 1986) that optimizes labeling on a voxel-by-voxel basis rather than for the entire data set. The algorithm is initialized by assigning each image voxel the label that yields the lowest energy according to the first term in Eq. [A8]. To achieve a MAP (i.e., to satisfy Eq. [A8]), each voxel is relabeled in each iteration step based on the least energy contribution from that particular voxel due to all possible labels. A new energy is then calculated. The algorithm terminates when the energy change is below a certain threshold. The pseudocode for the implemented ICM algorithm is

1. Initialize a labeling $\hat{L}_{current} \leftarrow \hat{L}_0$, for which the energy $E_{current} \leftarrow E(\hat{L})$ is a minimum.
2. Repeat:
   i. calculate the energy of each voxel due to the label l_1: $e_{small} \leftarrow e(l_1)$
   ii. determine the label for each voxel that contributes to the least energy; repeat for $l_i = l_2, ..., l_n$.
      a. check if $e(l_i) < e_{small}$
      b. accept this label as the new desired label: $l_{new} \leftarrow l_i$
   iii. update the new label for the set of all voxels, $\hat{L}_{new}$, and calculate the corresponding energy, $E_{new} \leftarrow E(\hat{L}_{new})$
   iv. calculate the change in energy: $\Delta E \leftarrow E_{new} - E_{current}$
   v. update the current energy, $E_{current} \leftarrow E_{new}$

   until $\Delta E < E_{threshold}$

In Step 1, a labeling for all voxels, $\hat{L}_0$, that corresponds to the least energy contribution due to all possible labels for each voxel, l_i, is assigned and its corresponding energy is calculated. In Steps 2.i and 2.ii, the voxels are assigned new labels, l_{new}, that yield the lowest energy according to Eq. [A8]. In Step 2.iii, the resulting new labeling, $\hat{L}_{new}$, for the entire three-dimensional volume is updated and its corresponding overall energy is calculated. Every time a new labeling is assigned, the change in energy is calculated and the loop continues until this change is below a predefined threshold, i.e., the exit criterion.

Acknowledgments

We gratefully acknowledge support from the U.S. National Science Foundation (NSF) under Grant no. EAR-0911242 and from the Arizona Agricultural Experiment Station (AAES).

References

Baveye, P.C., M. Laba, W. Otten, L. Bouckaert, P. Dello Sterpaio, R.R. Goswami, et al. 2010. Observer-dependent variability of the thresholding step in the quantitative analysis of soil images and X-ray microtomography data. Geoderma 157:51–63. doi:10.1016/j.geoderma.2010.03.015

Berthod, M., Z. Kato, S. Yu, and J. Zerubia. 1996. Bayesian image classification using Markov random fields. Image Vis. Comput. 14:285–295. doi:10.1016/0262-8856(95)01072-6

Besag, J.E. 1974. Spatial interaction and the statistical analysis of lattice systems. J. R. Stat. Soc. B 36:192–236.

Besag, J.E. 1986. On the statistical analysis of dirty pictures. J. R. Stat. Soc. B 48:259–302.

Blum, C., and A. Roli. 2003. Metaheuristics in combinatorial optimization: Overview and conceptual comparison. ACM Comput. Surv. 35:268–308. doi:10.1145/937503.937505

Chandra, R., L. Dagum, D. Kohr, D. Maydan, J. McDonald, and R. Menon. 2001. Parallel programming in OpenMP. Morgan Kaufmann, San Francisco.

Chantas, G., N.P. Galatsanos, R. Molina, and A.K. Katsaggelos. 2010. Variational Bayesian Image restoration with a product of spatially weighted total variation image priors. IEEE Trans. Image Process. 19:351–362. doi:10.1109/TIP.2009.2033398

Chuang, K.-S., H.-L. Tzeng, S. Chen, J. Wu, and T.-J. Chen. 2006. Fuzzy c-means clustering with spatial information for image segmentation. Comput. Med. Imaging Graph. 30:9–15. doi:10.1016/j.compmedimag.2005.10.001

Deng, H., and D.A. Clausi. 2005. Unsupervised segmentation of synthetic aperture radar sea ice imagery using a novel Markov random field model. IEEE Trans. Geosci. Remote Sens. 43:528–538.

Derin, H., H. Elliott, R. Cristi, and D. Geman. 1984. Bayes smoothing algorithms for segmentation of binary images modeled by Markov random fields. IEEE Trans. Pattern Anal. Mach. Intell. 6:707–720.

Elliot, T.R., and R.J. Heck. 2007. A comparison of 2D vs. 3D thresholding of X-ray CT imagery. Can. J. Soil Sci. 87:405–412. doi:10.4141/CJSS06017

Geman, S., and D. Geman. 1984. Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. IEEE Trans. Pattern Anal. Mach. Intell. 6:721–741. doi:10.1109/TPAMI.1984.4767596

Glover, F., and G.A. Kochenberger. 2003. Handbook of metaheuristics. Int. Ser. Oper. Res. Manage. Sci. Kluwer Acad. Publ., Norwell, MA.

Goldberg, D.E. 1989. Genetic algorithms in search, optimization and machine learning. Addison-Wesley Longman Publ. Co., Boston.

Held, K., E.R. Kops, B.J. Krause, W.M. Wells III, R. Kikinis, and H.W. Muller-Gartner. 1997. Markov random field segmentation of brain MR images. IEEE Trans. Med. Imaging 16:878–886. doi:10.1109/42.650883

Holland, J.H. 1975. Adaptation in natural and artificial systems. Univ. of Michigan Press, Ann Arbor.

Iassonov, P., T. Gebrenegus, and M. Tuller. 2009. Segmentation of X-ray CT images of porous materials: A crucial step for characterization and quantitative analysis of pore structures. Water Resour. Res. 45:W09415. doi:10.1029/2009WR008087

Iassonov, P., and M. Tuller. 2010. Application of image segmentation for correction of intensity bias in X-ray CT images. Vadose Zone J. 9:187–191. doi:10.2136/vzj2009.0042

Joshi, M.A. 2006. Digital image processing: An algorithmic approach. Prentice Hall of India, New Delhi.

Kaestner, A., E. Lehmann, and M. Stampanoni. 2008. Imaging and image processing in porous media research. Adv. Water Resour. 31:1174–1187. doi:10.1016/j.advwatres.2008.01.022

L'Ecuyer, P., R. Simard, E.J. Chen, and W.D. Kelton. 2002. An object-oriented random-number package with many long streams and substreams. Oper. Res. 50:1073–1075. doi:10.1287/opre.50.6.1073.358

Levada, A.L.M., N.D.A. Mascarenhas, and A. Tannus. 2010. A novel MAP-MRF approach for multispectral image contextual classification using combination of suboptimal iterative algorithms. Pattern Recognit. Lett. 31:1795–1808. doi:10.1016/j.patrec.2010.04.007

Li, S.Z. 2009. Markov random field modeling in image analysis. 3rd ed. Springer-Verlag, London.

Lindquist, W.B. 1999. 3DMA general users manual. Tech. Rep. SUNYSBAMS-99-20. State Univ. of New York, Stony Brook.

Kato, Z., J. Zerubia, and M. Berthod. 1992. Satellite image classification using a modified Metropolis dynamics. p. 573–576. In IEEE Int. Conf. on Acoustics, Speech, and Signal Processing, San Francisco. 23–26 Mar. 1992. Vol. 3. IEEE, New York.

Kindermann, R., and J.L. Snell. 1980. Markov random fields and their applications. Am. Math. Soc., Providence, RI.

Kirkpatrick, S., C.D. Gelatt, and M.P. Vecchi. 1983. Optimization by simulated annealing. Science 220:671–680.

Metropolis, N., A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, and E. Teller. 1953. Equation of state calculations by fast computing machines. J. Chem. Phys. 21:1087–1092.

Moussouris, J. 1974. Gibbs and Markov random systems with constraints. J. Stat. Phys. 10:11–33. doi:10.1007/BF01011714

NVIDIA Corp. 2007. NVIDIA CUDA: Compute Unified Device Architecture programming guide, version 1.1. Tech. Rep. NVIDIA Corp., Santa Clara, CA.

Paulsen, R.R., J.A. Baerentzen, and R. Larsen. 2010. Markov random field surface reconstruction. IEEE Trans. Vis. Comput. Graph. 16:636–646. doi:10.1109/TVCG.2009.208

Porter, M.L., and D. Wildenschild. 2010. Image analysis algorithms for estimating porous media multiphase flow variables from computed microtomography data: A validation study. Comput. Geosci. 14:15–30. doi:10.1007/s10596-009-9130-5

Porter, M.L., D. Wildenschild, G. Grant, and J.I. Gerhard. 2010. Measurement and prediction of the relationship between capillary pressure, saturation, and interfacial area in a NAPL–water–glass bead system. Water Resour. Res. 46:W08512. doi:10.1029/2009WR007786

Rivers, M.L. 2010. GSECARS tomography processing software. Available at cars9.uchicago.edu/software/idl/tomography.html (verified 4 Dec. 2011). Univ. of Chicago, Chicago.

Schlüter, S., U. Weller, and H.-J. Vogel. 2010. Segmentation of X-ray microtomography images of soil using gradient masks. Comput. Geosci. 36:1246–1251. doi:10.1016/j.cageo.2010.02.007

Sudderth, E.B., and W.T. Freeman. 2008. Signal and image processing with belief propagation. IEEE Signal Process. Mag. 25:114–121. doi:10.1109/MSP.2007.914235.

Vaz, C.M.P., I.C. de Maria, P.O. Lasso, and M. Tuller. 2011. Evaluation of an advanced benchtop micro-computed tomography system for quantifying porosities and pore-size distributions of two Brazilian Oxisols. Soil Sci. Soc. Am. J. 75:832–841. doi:10.2136/sssaj2010.0245

Wang, W., A.N. Kravchenko, A.J.M. Smucker, and M.L. Rivers. 2011. Comparison of image segmentation methods in simulated 2D and 3D microtomographic images of soil aggregates. Geoderma 162:231–241. doi:10.1016/j.geoderma.2011.01.006

Wildenschild, D., J.W. Hopmans, M.L. Rivers, and A.J.R. Kent. 2005. Quantitative analysis of flow processes in a sand using synchrotron-based X-ray microtomography. Vadose Zone J. 4:112–126. doi:10.2113/4.1.112

Wildenschild, D., J.W. Hopmans, C.M.P. Vaz, M.L. Rivers, D. Rikard, and B.S.B. Christensen. 2002. Using X-ray computed microtomography in hydrology: Systems, resolutions and limitations. J. Hydrol. 267:285–297. doi:10.1016/S0022-1694(02)00157-9

Yip, P.P.C., and Y.-H. Pao. 1995. Combinatorial optimization with use of guided evolutionary simulated annealing. IEEE Trans. Neural Netw. 6:290–295. doi:10.1109/72.363466