

Statistical Fusion of Two-Scale Images of Porous Media

Azadeh Mohebi^a, Paul Fieguth^a, Marios Ioannidis^b

^a*Department of Systems Design Engineering, University of Waterloo, Waterloo, Ontario, Canada*

^b*Department of Chemical Engineering, University of Waterloo, Waterloo, Ontario, Canada*

Abstract

The reconstruction of 3D architecture of void space in porous media is a challenging task, since porous media contain pore structures at multiple scales. We propose a statistical fusion framework for reconstructing high resolution porous media images from low resolution measurements. The proposed framework is based on a posterior sampling approach in which information obtained by low resolution MRI/NMR measurements are combined with prior models inferred from high resolution microscopic data, typically 2D. In this paper, we focus on two-scale reconstruction tasks in which the measurements resolve only the large scale structures, leaving the small scale to be inferred from the prior model.

Key words: Porous media reconstruction, Magnetic resonance imaging, Data fusion, Posterior sampling, Simulated annealing

1. Introduction

The presence of void structure at multiple length scales [19] poses a major challenge to understanding fluid transport in natural porous media. One example is mineral soils, the void space of which comprises a mixture of pore systems dominated by different pore-length scales in the *nm*- to *mm*-size range and aris-

ing from the very broad grain size distribution of sandy material and the large fraction of clay and silt typically present in soils [20]. Sedimentary rock formations, both sandstone and carbonate, which are important aquifers and hydrocarbon reservoirs, can also be highly heterogeneous. This is typically the case for carbonate rocks, where multiple pore systems, ranging in size from less than one *nm* to several *mm* [21], are the outcome of complex diagenetic processes taking place between the initial deposition and exploitation times. The presence of a very broad spectrum of pore length scales implies a certain degree of spatial organization [19], which manifests itself as heterogeneity over multiple length scales [18].

Three-dimensional images of the void structure are of fundamental value to understanding fluid transport in natural porous media. An ideal 3D image would be of sufficiently large size to capture heterogeneity over multiple length scales and of sufficiently high resolution to resolve the finest pore length scale of interest to transport. Unfortunately, despite significant advances [22, 23, 24] such images cannot be obtained using any single imaging modality presently available. The purpose of this contribution is to lay the ground for an alternative approach, namely the fusion of image data of diverse origin, resolution or dimensionality.

3D images of the internal structure of porous media can be obtained by stochastic reconstruction [3, 25, 7, 26], sampling processes by which *typical* samples of the porous material of interest are computationally generated. Depending on the absence or presence of measurements, stochastic reconstruction is a *prior* or *posterior* sampling process, respectively. Nearly all published porous media reconstruction studies concern prior sampling, in the sense that the reconstruction is based solely on a learned prior model. The prior model typically consists of one or more statistical functions (e.g. two-point correlation, chord-length

distribution, etc.) learned from 2D images of the pore space. In view of the ever increasing access to 3D imaging tools, the situation in which 3D tomographic or magnetic resonance imaging (MRI) data are available for a specific physical sample is more interesting. Although 3D tomographic [22] or MRI [21, 18] measurements fail to resolve fine scale pore structure, the fine scale *is*, in fact, partially constrained by the measurements. Thus, the possibility exists to infer high-resolution details from the pattern of low-resolution measurements. Within the framework of stochastic reconstruction, the process then is one of *posterior sampling*, aimed at fusing the information present in low-resolution 3D image data with a prior-model based on high-resolution 2D images.

A first attempt at fusing low resolution measurements with high resolution data has been recently made by Okabe and Blunt (2007). These authors have considered low-resolution 3D tomographic images of the pore space, in which only the larger pores are resolved, and employed stochastic reconstruction to resolve smaller-scale structure. In their method, small-scale structure which is unresolved in the 3D measurements is generated according to a prior model in a manner independent from the 3D measurements. There are two problems with this approach. One is that the measurements are not explicitly coupled to the prior model in computing a reconstruction. The other, more important one, is that only a small portion of the information provided by the low resolution 3D measurements is exploited, namely the overall sample porosity and the presence of void and solid domains. Low-resolution measurements, however, can be rich sources of information. X-ray computed tomography or MRI resolve not only the relatively larger pores, but also provide information on the local (voxel-scale) porosity of unresolved scales. In addition, MRI can provide 3D maps of parameters sensitive to the geometry and connectivity of unresolved scales [12].

For example, Pomerantz et al. (2008) have recently presented spatially-resolved measurements of the decay of transverse magnetization in a series of sedimentary rocks and analyzed them to obtain the spectrum of decay constants, T_2 , at the voxel scale. The physical interpretation of a T_2 spectrum as a distribution of pore surface-to-volume ratio, S/V , affords additional insight into the geometry (pore size) of unresolved length scales [27]. Such information can enrich the stochastic reconstruction process.

In this paper we propose a Bayesian statistical framework for explicitly fusing different types of low-resolution measurements with a high-resolution prior model. The paper is structured as follows. We begin by establishing the basis of image data fusion and modeling in terms of Gibbs random fields (GRF), briefly reviewing prior image models and paying particular attention to the contributions of measurements and prior model in the Bayesian framework. In a subsequent section, we show how stochastic reconstruction may be handled by simulated annealing as a posterior sampling process. We then evaluate statistical fusion of low-resolution measurements with a high-resolution prior model for a variety of real and synthetic images of porous media. The reconstruction results are shown to possess two-scale structures, consistent with the measured samples, while obeying the same statistical features encoded in the prior model. Relative to the method of Okabe and Blunt (2007), our reconstructions are closer to the physical sample, show fewer artifacts and are more tolerant of the measurement resolution.

2. Data Fusion and Modeling

Our proposed fusion method is based on a Bayesian framework, such that we wish to fuse low resolution measurements with a high resolution prior model.

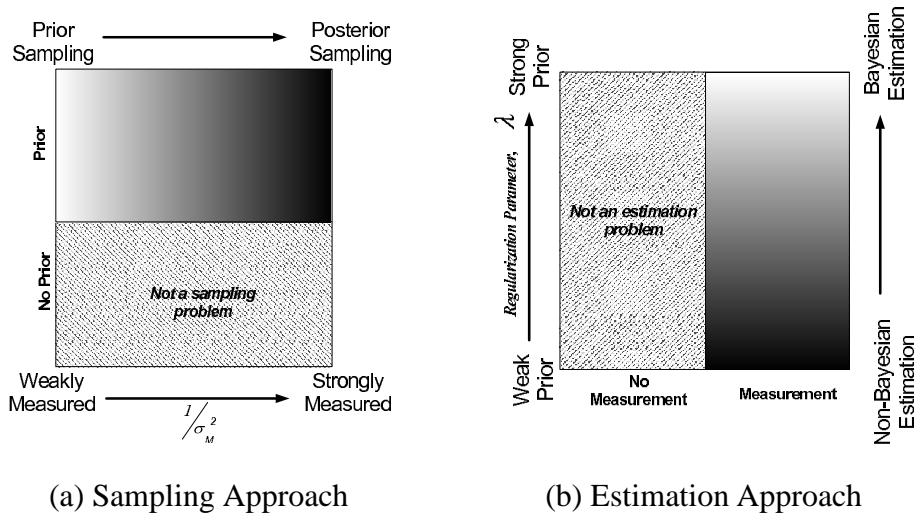


Figure 1: Sampling and estimation. In sampling (a) there is a continuum between prior and posterior sampling, depending on the measurement noise variance σ_M^2 , the degree to which the unknown random field is constrained by the measurements relative to the prior. In estimation (b) there is a continuum between Bayesian and non-Bayesian problems, depending on the regularization parameter λ , the degree to which a prior model is present to constrain the random field.

The measurements available for reconstruction may be (i) weakly constraining, (ii) partly constraining, or (iii) fully determining, a continuum illustrated in Figure 1(a). A great many problems in image processing fall into (iii), such that the measurements are of such a resolution and quality that no prior is needed, and the image is processed or reconstructed in the absence of any specific prior. At the other extreme (i) the measurements are so weak that one is essentially randomly sampling from a prior model. Our interest lies in between, such that something of a delicate balance is required between the assertions of measurements and prior model. This is the case with MRI or tomographic data of porous media in which relatively large pores are clearly resolved but smaller scale porosity appears in the images as different shades of gray.

In principle, given measurements and a prior model, reconstruction can proceed as either a sampling or estimation problem, as illustrated in Figure 1. Since our measurements fail to resolve many of the fine scale details and at the same time fine structure must be present in the stochastic reconstruction, a sampling approach is required [8]. Figure 2 shows how the posterior sampling approach combines the measurements with the high resolution samples to generate reconstructed porous media samples.

We propose to characterize the chaotic and complex morphology of porous media with the widely-used Gibbs Random Field (GRF) [15], defined as

$$p(Z) = \frac{e^{-H(Z)/T}}{\mathcal{Z}} \quad (1)$$

where $H(Z)$ is the *energy function*, and \mathcal{Z} is a normalization factor, termed the partition function.

Given measurements M , the posterior probability distribution is

$$p(Z|M) = \frac{e^{-H(Z|M)/T}}{\mathcal{Z}_M} \quad (2)$$

$$H(Z|M) = H(Z) + \alpha G(Z; M) \quad (3)$$

where α balances the contributions of prior and measurements. G constrains Z relative to the measurements,

$$G(Z; M) = \|f(Z) - M\| \quad (4)$$

where the forward model $f(\cdot)$ describes the physics or mathematics of the measurement process. We use the ℓ^2 -norm for $\|\cdot\|$, however the Gibbs formulation allows for other norm choices as well.

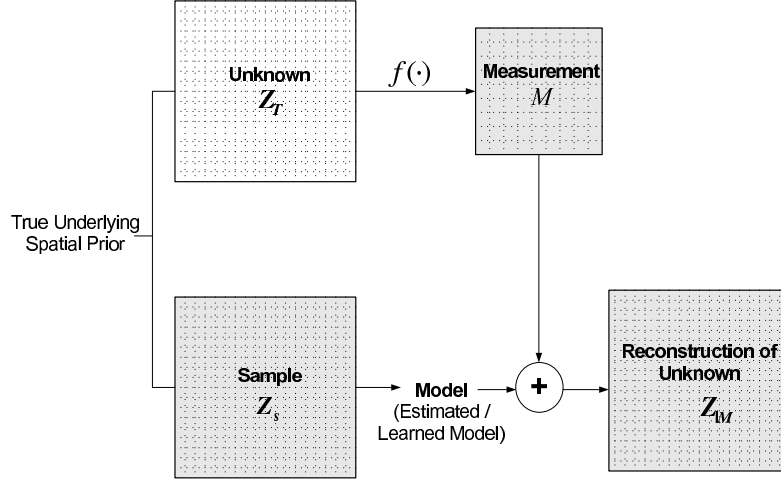


Figure 2: Posterior Sampling: The sample image Z_s is used to learn the prior model. A low resolution measurement M is taken from the unknown image Z_T . The unknown image Z_T and the sample image Z_s are assumed to obey the same statistics.

2.1. Prior Models

A variety of energy functions $H(\cdot)$ used to model porous media have been reviewed by [1]. These include autocorrelation, chord length distribution and lineal path functions, to name a few. For example the chordlength distribution function [1] $C^i(\ell)$ is defined to be the probability of finding a chord with length ℓ in phase i . For a two-phase porous medium, chordlength distributions can be defined for either phase and for chords at different orientations. It is common to limit the orientation to the horizontal and vertical directions, in which case chordlength energy in phase i is

$$H_c^i(Z) = \sum_{k=0}^L \|\bar{C}_h^i(k) - C_h(k, Z)\| + \|\bar{C}_v^i(k) - C_v(k, Z)\| \quad (5)$$

where \bar{C}_h^i and \bar{C}_v^i are the learned model for the horizontal and vertical chords, respectively, and C_h^i and C_v^i are the chordlength distributions of a simulated random field Z .

The histogram distribution function [2] is non-parametric, keeping the entire joint probability distribution of a local set of pixels within a neighborhood. Choosing eight adjacent pixels as the neighborhood structure leads to a non-parametric model containing a histogram of $n = 2^9$ probabilities. The histogram represent the probability of having each of 512 configurations in the image. The histogram energy function is defined as

$$H_h(Z) = \sum_{k=1}^{2^9} \left(\frac{\|\bar{h}(k) - h(k, Z)\|}{v(k) + \epsilon} \right) \quad (6)$$


where \bar{h} is the learned histogram distribution. and $h(Z)$ is the observed distribution for a given, simulated Z . The term v is the variance for the respective histogram entry, to account for sample variation. A small constant ϵ is introduced to avoid divisions by zero, especially in the comparatively common case of unobserved configurations k corresponding to $\bar{h}(k) = 0$.

In this paper we use a combined chordlength and histogram model, where the chordlength model asserts nonlocal statistics while the histogram model gives a detailed description of local structure.

2.2. Measurements as Constraints

A typical low resolution measurement of porous media obtained by computed tomography or MRI measures the *local* porosity of a given medium. The low resolution measurement, M_p , measures the porosity of a given material within a pixel in 2D or a voxel in 3D. Therefore, each datum in the porosity measurement corresponds to the void fraction of a given block of pixels of the underlying high resolution field Z .

Figure 3 plots porosity measurements at three different resolutions for two fields, the first synthetic and the second from sintered glass beads. The decimation parameter d measures the reduction in resolution from Z to M , such that



(a) Both scales resolved, $d = 5$ (b) Single scale resolved, $d = 15$ (c) No scales resolved, $d = 50$

Figure 3: Porosity measurements at different resolutions (with decimation factor d). Particularly in (c) nearly all information of the image structure has been lost, and the distinction of the two phases in the image is not inferrable from the measurements.

each single measurement measures $d \times d$ elements in two-dimensional Z . In (a) the measurements resolve all aspects of the field, leaving little to do for a prior model, whereas in (c) the measurements are relatively uniformly grey, meaning that any reconstructed sample is unlikely to bear a resemblance to the measured physical sample. Of greatest interest in this paper are the measurements of (b), in which large-scale pores are resolved, but fine-scale pores are not, leading to a balancing of assertions and contributions between the prior and measurements.

Other measurements can be defined which can provide complementary information on fine-scale structures. Consider, for example, a spatially-resolved measurement of the decay of transverse nuclear magnetization, $M_z(t)$ [18]. This measurement is generally described in terms of a multi-exponential distribution

of apparent relaxation times T_{2i}

$$\frac{M_z(t)}{M_0} = \sum f_i \exp(-t/T_{2i}) \quad (7)$$

where f_i is the volume fraction of fluid relaxing at a rate $1/T_{2i}$. A number of processes can potentially influence magnetization decay in porous media, including bulk and surface-enhanced relaxation and relaxation due to diffusion in the internal and external gradients. The effects of relaxation due to diffusion can be minimized by a proper choice of the pulse sequence and the effect of bulk relaxation can be accounted for, so that the distribution of surface-enhanced relaxation times T_{2s} may be extracted from the distribution of apparent relaxation times. A T_{2s} -distribution may be interpreted as a volume-weighted pore size distribution. That is, each relaxation rate $1/T_{2s}$ corresponds to a different pore length scale ℓ , where ℓ is the ratio of pore volume to pore surface area, V_p/S , of a region of pore space where uniform magnetization is maintained by molecular diffusion over measurement times [27]. MRI is thus able to provide a measurement M_d , reflecting information on the average surface-to-volume ratio within a pixel in 2D or a voxel in 3D. Each datum in measurement M_d reflects the average pore size of a given block of pixels of the underlying high resolution field Z . Such a measurement provides valuable clues with respect to unresolved structures and geometry, as illustrated in Figure 4.

Including both measurements leads to a modification of the posterior energy of Eq. (4) as

$$G(Z; M) = G(Z; M_d, M_p) = \|f_d(Z) - M_d\| + \gamma \|f_p(Z) - M_p\| \quad (8)$$

where $f_d(\cdot)$ and $f_p(\cdot)$ are the forward models for surface-to-volume ratio and porosity measurements, respectively, and γ is a weighting parameter between the two constraints.




Figure 4: Surface-to-volume ratio measurements at three resolutions. The measurement is a function of pore structure, even at unresolved scales. In sharp contrast to Figure 3, the delineation of the two regions in the top row is clear at even very low resolutions.

The forward model for obtaining the local porosity measurement, f_p is defined based on downsampling the simulated Z . Therefore a datum m_i in $f_p(Z)$ is defined as

$$m_i = \sum_{k=1}^{d^2} \frac{1}{d^2} B_i(k) \quad (9)$$

where d is the decimation factor and B_i is a $d \times d$ block in Z corresponding to the i th data in the measurement.

Defining f_d requires having surface and volumes, equivalently perimeter and area in 2D, of pores for a datum m_i in the measurement. For a block of pixels in Z , B_i , the perimeter is defined based on the fraction of pixels which are at the perimeter of pores, and the area is simple the local porosity measurement.

3. Posterior Sampling

Sampling from the Gibbs probability distribution is straightforward, in principle, by applying the Gibbs sampler [13, 14] to the posterior energy Eq. (2):

$$p(Z|M) = \frac{e^{-H(Z|M)/T}}{\mathcal{Z}_M} \quad (10)$$

In practice such sampling is challenging, since we do not know what value of temperature T to select for the sampler. This uncertainty stems from how energy models $H(Z|M)$ are developed for Z — we really don’t know any “true” model, rather we infer plausible constraints from sample data, such that implausible $Z|M$ is associated with a larger energy H , and very likely $Z|M$ with a correspondingly small value. Because of the highly qualitative notions of plausibility and implausibility, it is difficult (impossible) to select a value of T which leads to an accurate posterior distribution $p(Z|M)$.

Instead, the posterior sampling problem is most effectively understood in the context of simulated annealing [13]. Simulated annealing runs a Gibbs sampler, initially at a high temperature where the image contains purely random structures, then at progressively lower temperature until the system reaches a converged state. In other words, the annealing procedure generates a Markov chain which converges in distribution to the uniform measure over the minimal energy configurations.

If we view the measurements as noiseless and the prior energy as a set of empirical hard constraints, then the set of minimal energy configurations becomes

$$\{Z \mid H(Z) = 0 \text{ and } G(Z; M) = 0\} \quad (11)$$

That is, the posterior sample is a random Z chosen from the combined null-spaces of the prior model H and measurement model G .

In practice, there are multiple reasons why this hard-constraint perspective becomes impractical:

1. It is possible for the null spaces of H and G to be non-intersecting, meaning that *no* solution exists for Z . In other words, there might be inconsistencies between the information provided by the prior model and the measurements, such that a valid structure in one may be forbidden in the other.
2. An energy function which contains terms which either permit or forbid hypothesized states leads to an energy map with very deep local minima, which present difficulties to simulated annealing. Continuous energy functions are more likely to lead to robust convergence.
3. Given a finitely sized sample, it is not possible to infer *rigid* rules and constraints regarding the prior model. At best we can infer that certain behaviors are more or less likely to occur, again leading to a continuous definition for the prior energy H .

If the measurements are treated as exact, then our posterior sampling problem is an example of annealing subject to hard constraints, a problem which has previously been investigated [13]. We now seek a random sample from $p(Z)$, subject to precisely matching the measurements:

$$\{Z \sim p(Z) \mid G(Z; M) = 0\}. \quad (12)$$

Since the hard constraint is not easily asserted, in practice this is most easily accomplished within the context of annealing by asserting the hard constraint to an ever increasing degree, annealing over

$$H(Z|M) = H(Z) + \alpha G(Z; M) \quad (13)$$

such that $\alpha \rightarrow \infty$ as $T \rightarrow 0$. The result of this annealing is to produce a posterior sample, forced to be consistent with the measurements, and randomly sampled from the prior within the constrained space.

In practice we have used an exponential cooling schedule for T . This process starts with high temperature, large T , and small α , and continues by decreasing T very slowly with constant or slowly increasing α . The sampling process starts with a purely random binary image, without any special initialization or porosity assertion, since the measurement in the energy function will itself constrain on the local porosity during the annealing process. When the measurement constraints are satisfied, the temperature is fixed at $T_c \neq 0$ and the Gibbs sampler continues sampling from the constrained space in Eq. (12), at non-zero temperature. At the end the posterior samples are consistent with prior model while being faithful to the measurements.

In the more usual event of imprecise measurements, the setting of the relative weight α as a function of annealing iteration is less clear and remains an open problem.

4. Results and Evaluation

We have applied constrained sampling on various data shown in Figure 5 (a): the small-large circle toy problem, a sample of real vuggy carbonate rock, and two physical models of vuggy carbonate rock made of sintered glass spheres, recently studied by Padley et al. (2007). Based on repeated tests, all examples are based on a chordlength prior was selected, with the exception of the first sintered spheres, for which a histogram distribution acts as the prior model. In all cases an exponential annealing schedule was used, as follows

$$T_k = T_0 b^{(k-1)} \tag{14}$$


where k indicates the iteration and T_0 is the starting temperature. The parameter γ in Eq. (8) was set to make equal the initial energies of the porosity and surface-to-volume ratio measurements, and the initial value of α in Eq. (13) was set to make equal the initial energies of the prior and measurement terms.

A first set of samples, generated by the proposed constrained sampling approach, are shown in Figure 5. In row (d) of this figure we can see a stunning enhancement of the low-resolution measurements of rows (b) and (c). Indeed, (d) reproduces many of the finer-scale features seen in the original images of (a), but which are completely lost in the measurements (b). Although the prior model does not have any specific preference for the two different structures represented in the small-large circle data, the reconstructed sample *does* contain structures at two distinguishable scales due to using the information provided by the measurements. The samples are not perfect, of course, and are particularly limited by the sophistication and quality of the prior model. In particular, our use of a single, stationary prior model has some difficulty in reproducing fine details in the nonstationary fields of the leftmost and rightmost columns in Figure 5.

Figure 6 shows three posterior samples of sintered glass spheres, all three drawn from the same prior model and the same measurements. That the three samples are different, despite being drawn from identical prior and measurements, serves to emphasize the *random* posterior sampling nature of our approach. Features which are resolved by the low resolution measurements, such as large pores, are preserved the same way in each image, whereas unresolved features are randomly sampled.

To evaluate our results we begin with Figure 7, which shows the correlation between the original (original) and reconstructed data. Because a pixel in the middle of a large pore or solid is likely to be the same in the original and recon-


Small-large circle toy problem Real carbonate rock Sintered glass sphere(1) Sintered glass sphere (2)




(a) Original media



(b) Low resolution porosity measurements, downsampled by $d = 15$ from (a).



(c) Low resolution surface-to-volume ratio measurements



(d) Reconstruction, using both measurements in (b) and (c)

Figure 5: Artificially reconstructed samples (bottom), using posterior sampling, based on the low-resolution measurements in (b) and (c). The improvement in detail in (d) relative to (b) and (c) is striking.




Figure 6: Three posterior samples: observe that the large, resolved structures (such as the pore, top-right and bottom) remain unchanged, whereas unresolved details (fine-scale structure) are randomly synthesized. Thus we are constructing multiple samples, all representative of a given measured medium.

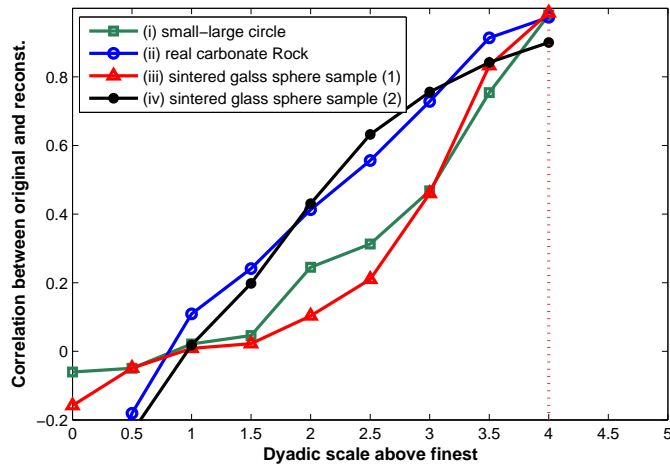


Figure 7: The correlation between the reconstruction results in Figure 5(d) and their corresponding original images. The dotted line shows the measured scale. It is clear that the reconstruction is correlated with the true field one or two scales finer than measured. That is, the proposed posterior sampling approach *does* add value below the measured scale; the fine-scale details are not just *random*, rather they are consistent with the measurements, where constrained by the measurements, and consistent only with the prior model where not constrained by the measurements.

structed images, as opposed to a pixel on a pore or solid boundary, the correlation is computed as a function of scale, where the scale of each pixel is defined as the number of scales over which a pixel's value is unchanged under repeated decimation. The resulting figure reflects the degree to which the results in Figure 5 (c) are consistent with the original data. As seen in Figure 7, the correlation is significant even below the measured scale, shown as a vertical dotted-line, emphasizing that the structures in the reconstructed results *below* the measurement scale remain consistent with the original data.

We have also compared our proposed method with another method described by Okabe and Blunt (2007), which proposes a modified form of posterior sampling, in which areas where measurements are purely white or black (i.e., pore or solid) are preserved, with all of the remaining values filled in by *prior* sampling. Because the fraction of black or white pixels in the measurements is clearly a decreasing function of the decimation factor d , as shown in Figure 8, the performance of the method of Okabe and Blunt (2007) is likely to suffer as d increases. Indeed, the results in Figure 9 confirm this conjecture, and furthermore show the strength of the proposed posterior sampling approach. Figure 9(c) suffers from blocky structures and pores completely disappearing as they fail to be explicitly resolved in the poorer measurements. The key to our results is that a given feature is not either perfectly resolved or fully unresolved, rather the degree to which the measurements resolve a feature lies on a continuum. Even grey, low-resolution measurements offer *some* constraint on the reconstruction, and it is the *interaction* between measurements and prior model which provides the robust results, specifically with respect to the middle pore, in Figure 9(d).

Next, Figure 10 shows the correlation, computed as in Figure 7, between the reconstructed samples and original data for three different approaches:

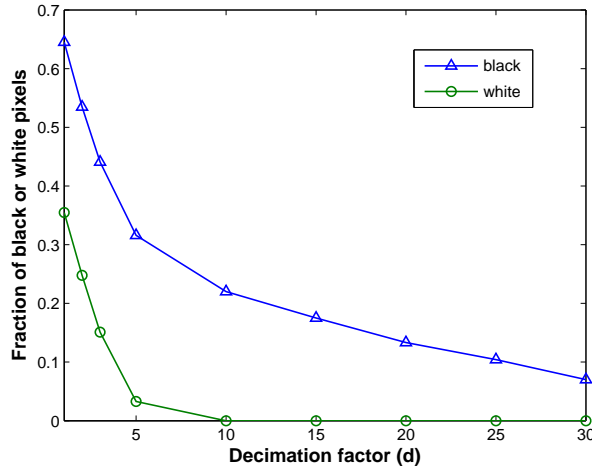



Figure 8: Fraction of purely black or purely white pixels in the measurements as a function of downsampling for the sintered glass spheres.

1. Random sampling (no prior) constrained by the measurements,
2. Prior sampling constrained on pure black and white pixels [11], and
3. Our proposed posterior sampling.

In the first method no prior is involved and the results are generated randomly so that they would be only consistent with the measurements. In the second method proposed by Okabe and Blunt (2007), the pure solid and pore areas are unchanged in the reconstructed image, while the rest are generated through a prior sampling approach. As is expected, using only measurements (first method) or using uncoupled measurements and prior (second method) lead to reconstructions in which little or no relevant detail is created at scales finer than the measured one, whereas our coupled, measurement-prior posterior sampling approach leads to positively correlated details.

A similar conclusion can be reached by examining the overall reconstruction accuracy, measured in terms of Mean Squared Error (MSE) between the original






(a) The original data




Decimation $d = 10$

Decimation $d = 15$




Decimation $d = 20$



(b) Local porosity measurement at different resolutions

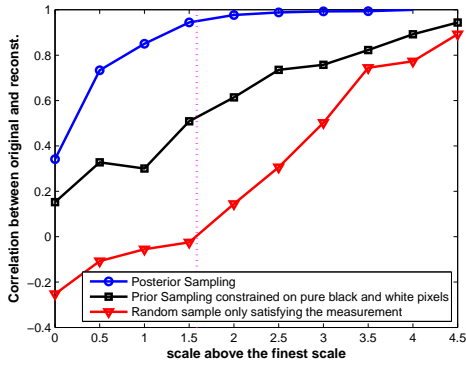


(c) Reconstruction, using prior sampling constrained on pure black and white pixels in the measurements [11]




(d) Reconstruction, using the proposed posterior sampling with porosity and surface-to-volume ratio measurements

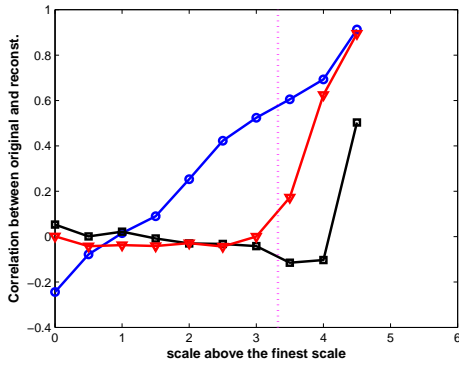
Figure 9: Reconstructed artificial samples. Because the method of (Okabe and Blunt, 2007) relies on the presence of purely white and black measurements, as d increases the reconstructions suffer from blocky artifacts and lost pores (see middle of each image). Posterior sampling, on the other hand, is able to use shades of grey as constraints on the random field, and therefore is able to tolerate reduced resolution more gracefully.




(a) $d=3$



(b) $d=5$



(c) $d=10$



(d) $d=30$

Figure 10: Correlation between reconstructed and original data (for sintered glass spheres) comparing our proposed approach (posterior sampling) with two other methods, as a function of scale, for four different values of the decimation parameter d . The dotted line shows the measured scale in each case. The proposed method is more successful in reconstructing structures finer than the measured scale, with the distance between the blue and black curves essentially representing the improvement of our proposed method over that of Okabe and Blunt (2007).

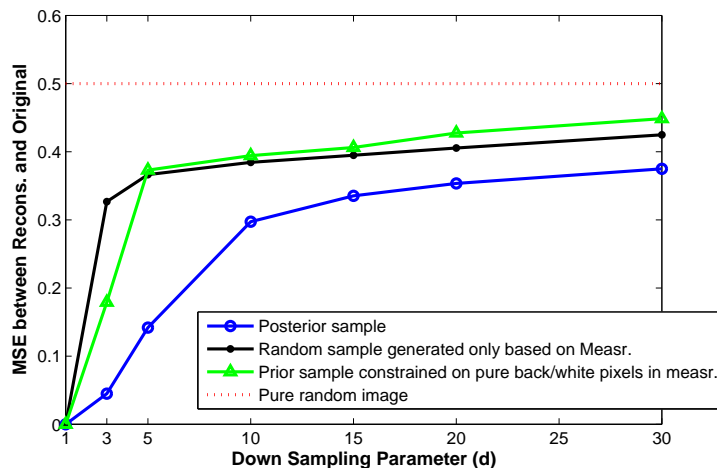


Figure 11: Mean-Square Error (MSE) between the reconstructed and true fields as a function of decimation factor d for sintered glass spheres. The four lines show MSE corresponding to a purely random field, a prior-free reconstruction based on measurements only the prior sample constrained on pure black and white pixels in the measurement, and the posterior sample which we produce.

and reconstructed samples. Figure 11 plots the MSE for the three different methods described above as a function of decimation factor d . Compared with the two other approaches, posterior sampling produces results most consistent with the original sample.

Finally, in Figure 12 we illustrate the application of our proposed approach to real MRI measurements. The prior model is learned from accurate, high-resolution images, shown in Figure 12(a). However, in contrast with previous experiments, in which we synthesized the measurements from high resolution data, in this case the measurements are actual MRI data [21], taken from a different, but statistically consistent, sample. The resulting reconstructed results in Figure 12(c) are much higher in resolution than the measurements in (b). We can


see that the large scale structures in both examples are preserved, and that unresolved small scale structures are mostly consistent with the underlying statistics of (a). Since we do not have the measurement corresponding to S/V for the real samples at this time, the reconstruction is only based on one type of measurement — the porosity measurements shown in (b). We would have obtained a better enhancement in resolution and more consistent structures, if the S/V measurement was also considered in the reconstruction process. Moreover, due to the significant noise in the measurements we also expect, and see, noisy structures in the reconstructed samples primarily, we hypothesize, because the chosen prior models are insufficiently discriminating.

5. Conclusion and Future Work



In this paper we proposed a statistical fusion approach based on posterior sampling for two-scale porous media reconstruction, in particular, with problems having one resolved and one unresolved scale. The key to our approach is the simultaneous assertion of prior and measurement constraints, which leads to superior reconstructions, possessing useful details at scales finer than that of the measurements.

We have identified a limitation of computational complexity, due to the slow convergence of simulated annealing, and a limitation of modeling sophistication, due to the assertion of a prior model on a single scale. Our future work seeks to address both of these issues by methods of hierarchical sampling.


An extension of the proposed approach to three dimensions should be possible in principle, however again the method is likely to be limited by computational complexity, unless based on a hierarchical approach — essentially along the lines of what has been proposed in [2]. Consequently our future research fo-



(a) 2D high resolution samples (512×512)



(b) Real 2D local porosity measurements (32×32) [21]



(c) Reconstruction, using porosity measurement (515×512)

Figure 12: Reconstruction of two types of porous media using *real* 2D porosity measurements. The prior is learned from a 2D high resolution sample (a), and a single frame of 3D MRI low-resolution sample, shown in (b), is used as the measurement. The posterior samples in (c) are the reconstructed samples of (b), 16 times higher in resolution. The samples in (a) and (b) are not identical, but are assumed to obey the same statistics.

cuses on the development of superior prior models, based on hierarchical modeling, both to enable scale-dependent models, and to lead to an implementation for 3D reconstruction.

References

- [1] Torquato S. *Random heterogeneous materials: Microstructure and macroscopic properties*. Springer-Verlag, 2002.
- [2] Alexander S., Fieguth P., Ioannidis M., Vrscay E. Hierarchical Annealing for Synthesis of Binary Porous Media Images. *accepted in Mathematical Geoscience*, 2008.
- [3] Adler P. M. *porous media, geometry and transports*. Butterworth-Heinemann series in chemical engineering, 1992.
- [4] Yeong C. L. Y. and Torquato S. Reconstructing porous media. *Physical Review E*, 57(1):495-506, 1998.
- [5] Yeong C. L. Y. and Torquato S. Reconstructing random media II. Three-dimensional media from two-dimensional cuts. *Physical Review E*, 58(1):224-233, 1998.
- [6] Ioannidis M. A. and Chatzis I. On the geometry and topology of 3D stochastic porous media. *Journal of Colloid Interface Science*, 229:323-334, 2000.
- [7] Manwart C., Torquato S. and Hilfer R. Stochastic reconstruction of sandstones. *Physical Review E*, 62(1):893-899, 2000.
- [8] Fieguth P. Hierarchical posterior sampling for images and random fields. *IEEE International Conference on Image Processing*, 2003.
- [9] Mohebi A. and Fieguth P. Posterior Sampling of Scientific Images. *LNCS*, 4141:365-376, 2006.
- [10] Zhao X., Yao J. and Yi Y. A new stochastic method of reconstructing porous media. *Transport in Porous Mediaing*,
- [11] Okabe H. and Blunt M. J. Pore space reconstruction of vuggy carbonates using microtomography and multiple-point statistics. *Water Resource Research*, 43, 2007.69(1):1-11 , 2007.

- [12] Stapf S. and Han S. *NMR imaging in chemical engineering*. Wiley-VCH, 2006.
- [13] Geman D. Random fields and inverse problems in imaging. In P. L. Hennequin, editors. Springer-Verlag, Lecture Notes in Mathematics, 1427, 1990.
- [14] Geman S. and Geman D. Stochastic relaxation, Gibbs Distribution, and the Bayesian Restoration of Images. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 6(6), 1984.
- [15] Gelfand S. B. and Mitter S. K. *Markov random fields, theory and applications*. Academic Press, 1993.
- [16] Olayinka S. and Ioannidis M. Time-dependent diffusion and surface-enhanced relaxation in stochastic replicas of porous rock. *Transport in Porous Media*, 54:273-295, 2004.
- [17] Winkler G. *Image analysis, Random Fields, and Markov Chain Monte Carlo Methods*. Springer-Verlag, 2003.
- [18] Pomerantz A. E., Sigmund E. E., Song Y. Q. Spatial Heterogeneity Length Scale in Carbonate Rock. *Applied Magnetic Resonance*, 32:221-231, 2007.
- [19] Radlinski A., Ioannidis M., Hinde A., Hainbuchner M., Baron M., Rauch H., and Kline S. Angstromto- millimeter characterization of sedimentary rock microstructure *Journal of Colloid and Interface Science*, 274(2):607612, 2004.
- [20] Tsakiroglou, A. and Ioannidis, M. Dual-porosity modelling of the pore structure and transport properties of a contaminated soil *European Journal of Soil Science*, 59:744-761, 2008.
- [21] Padhy A. and Ioannidis M. Induced pressure gradients due to entrance and exit effects in electroosmotically driven flows through nanopores within the continuum regime *Journal of Physics D-Applied Physics*, 41(6), 2008.
- [22] Sakellariou A, Arns CH, Sheppard A., et al. Developing a virtual materials laboratory *Materials Today*, 10(12):44-51, 2007.
- [23] Mobus G., Inkson B. Nanoscale tomography in materials science *Materials Today*, 10(12):18-25, 2007.
- [24] Kaestner A., Lehmann E., Stampanoni M. Imaging and image processing in porous media research *Advances in Water Resources*, 31(9):1174-1187, 2008.
- [25] Roberts A. Statistical reconstruction of three-dimensional porous media from

- two-dimensional images *Physical Review E*, 56(3):3203-3212, 1997.
- [26] Talukdar M., Torsaeter O. and Ioannidis M. Stochastic reconstruction of particulate media from two-dimensional images *Journal of Colloid and Interface Science*, 248(2):419-428, 2002.
- [27] Olayinka S. and Ioannidis M. Time-dependent diffusion and surface-enhanced relaxation in stochastic replicas of porous rock *Transport in Porous Media*, 54(3):273-295, 2004.