Posterior Sampling of Scientific Images

Azadeh Mohebi and Paul Fieguth

Department of Systems Design Engineering, University of Waterloo
Waterloo, Ontario, Canada, N2L 3G1

Abstract. Scientific image processing involves a variety of problems including image modelling, reconstruction, and synthesis. We are collaborating on an imaging problem in porous media, studied in-situ in an imaging MRI in which it is imperative to infer aspects of the porous sample at scales unresolved by the MRI. In this paper we develop an MCMC approach to resolution enhancement, where a low-resolution measurement is fused with a statistical model derived from a high-resolution image. Our approach is different from registration/super-resolution methods, in that the high and low resolution images are treated only as being governed by the same spatial statistics, rather than actually representing the same identical sample.

1 Introduction

Scientific imaging plays a significant role in research, especially with the availability of sophisticated imaging tools, including magnetic resonance imaging, scanning electron microscopy, confocal microscopy, computer aided X-ray tomography, and ultrasound, to name only a few. Because of the significant research funding and public interest in medical imaging and remote sensing, these aspects of scientific imaging have seen considerable attention and success.

However there is an enormous variety of imaging problems outside of medicine and remote sensing, where we would argue the current image processing practice to be relatively rudimentary, and where substantial contributions remain to be made. One such area is that of porous media [1] — the science of water-porous materials such as cement, concrete, cartilage, bone, wood, and soil, with corresponding significance in the construction, medical, and environmental industries.

The research in this paper is predicated on the quandary of imaging such porous media. High-resolution two-dimensional images of a porous surface can be produced, however viewing the interior of a sample requires cutting, polishing, and exposure to air, all of which may alter the sample. In-situ three-dimensional images of a medium can be measured using an imaging MRI, however the spatial resolution is very limited, such that only the largest pores are resolved.

Our long-term objective is the fusion of data from multiple imaging modalities to produce high-resolution images of porous media; for example, the fusion

A. Campilho and M. Kamel (Eds.): ICIAR 2006, LNCS 4141, pp. 339–350, 2006.

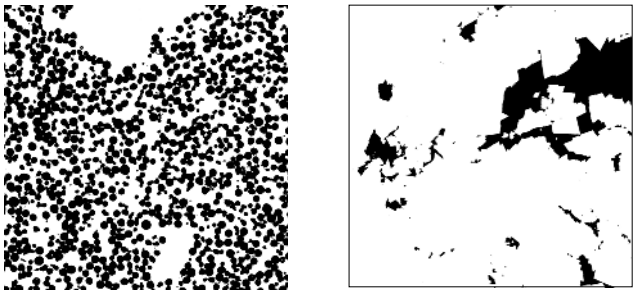

(a) Sintered Glass Spheres (b) Berea Sandstone

Fig. 1. Two examples of high-resolution 2D slices of irregular porous media

of high-resolution 2D and low-resolution 3D measurements. As it will never be feasible to acquire hundreds or thousands of 2D slices from a single sample, the given 2D measurements will image only an infinitesimal fraction of the 3D sample, so we are *not* interested in a 2D-3D co-registration problem, as is common in medical imaging and remote sensing. Instead, we view the 2D image as *characterizing* the sample, such that we infer a statistical model from the image, and fuse this model with the 3D data set to infer details at a higher resolution.

As an initial exploration of the above long-term research objective, this paper explores the statistical fusion of low and high-resolution data sets, limiting our attention here to two-dimensional random fields for simplicity.

2 Bayesian Image Analysis

Most porous media [1], such as those shown in Fig. 1, possess clear spatial patterns and relationships which characterize the medium. While simple image prior models can be obtained from image statistics such as spatial variances and correlation functions, such models are particularly poor for discrete-state (pore/solid) problems, in which case Gibbs random-field models are widely used [8],[5].

As the generation of a high resolution realization from a low resolution observation is highly ill-posed, some sort of regularization constraint is needed. In a non-Bayesian case (Fig. 2 (I,II)), multiple input images can be combined using methods of super resolution and data-fusion [4]. However, in our scientific imaging application in porous media we do not have the luxury of multiple images, although we may have available high-resolution measurements from statistically-equivalent samples (Fig. 2 (III)). Therefore, we consider a Bayesian approach in which the prior model can be characterized according to a high resolution 2D image and a low resolution 2D image or 3D volume which is the measurements characterizing a sample from which we wish to infer high resolution details.

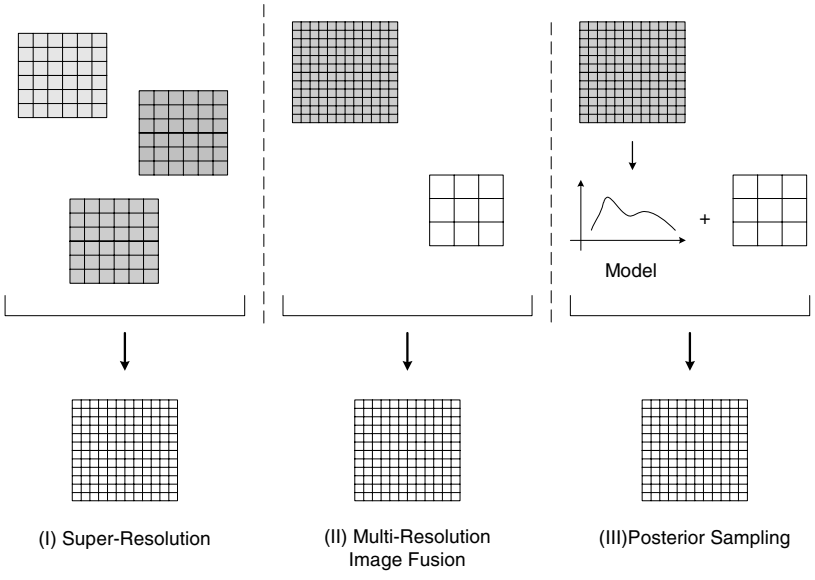

Fig. 2. Different prospectives in data fusion: We may be registering and fusing multiple images (I), fusing data across resolution (II), or statistical fusion through a model (III), as we propose in this paper. Although (II) and (III) appear superficially similar, in (II) the two data sets correspond to the identical underlying image, whereas in (III) the data are only assumed to obey the same statistics.

3 Posterior Sampling

There are two common objectives associated with Bayesian random-field models [9], [10]:

1. To generate random realizations consistent with the prior statistics of the model, referred to as random synthesis or prior sampling.
2. To solve for the optimal image based on measurements and a prior model, known as as estimation.

Prior sampling is not a function of measurements, and is therefore of limited use in settings where we wish to enhance a low-resolution data set, whereas estimation produce only those image features or structures which are inferable from the measurements, and therefore of limited utility in porous media which exhibit behavior over a wide range of scales. Instead, we propose to do posterior sampling (Fig. 3) — the synthesis of random images simultaneously obeying both the measurements and the prior models, producing results similar to estimates in densely-measured area, and producing a random synthesis in those area not constrained by measured values, thereby creating a high resolution result containing structure on a variety of scales.

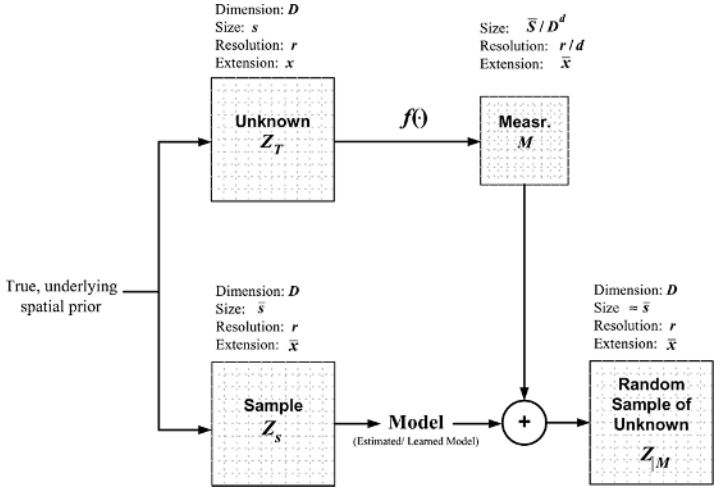

Fig. 3. Posterior Sampling. The sample image Z_s is used for learning/estimating the prior model. An irregular and incomplete measurement from the unknown image Z_T is available as M which is considered in the prior model. The posterior sample, unknown image and the sample image obey the same statistics.

3.1 Problem Formulation

To formulate the problem, we have used Gibbs Random Fields (GRFs) theory. GRFs are lattice models, used to quantify the spatial interactions of observed values at the nodes of a grid and to compute a probability for any configuration of that grid [12] [3]. GRFs were originally used in statistical physics to study the thermodynamic characteristics of interacting neighboring particles in a system [12]. Any Gibbs probability distribution takes the form

$$p(Z) = \frac{e^{-\beta H(Z)}}{\mathcal{Z}} \tag{1}$$

where $H(\cdot)$ is an energy function:

$$H(Z) = \sum_{c \in \mathcal{C}} V_c(Z) \tag{2}$$

written as a function of interactions over a local clique $\mathcal{C}$, where $\mathcal{Z}$ is a normalization factor, the *partition function*, and $\beta = 1/T$ is related to the temperature T. As the joint prior probability $p(Z)$ is strictly a function of H, the energy H implicitly encodes all of the characteristics of the random field. In order to solve the posterior sampling problem we will need to sample from the posterior probability [10]:

$$p(Z|M) = \frac{e^{-\beta H(Z|M)}}{\mathcal{Z}} \tag{3}$$

where

$$H(Z|M) = H(Z) + \alpha||M - f(Z)|| \tag{4}$$

for some norm $||\cdot||$, where $f(\cdot)$ is the forward model, describing how measurements are derived from unknowns,

$$M = f(Z). \tag{5}$$

Two questions remain: the selection of H, and the balancing of α and β in producing a random sample. The selection of H entails identifying the prior model and the measurement model, which are $H(Z)$ and $||M - f(Z)||$ respectively, shown in Eq. (4).

3.2 Prior Models

We consider two types of prior model:

1. *Ising model* [12]: Each site can take two possible values and a first-order neighborhood structure is considered for each site, i.e. for every site $z_{i,j}$ in the field, sites $z_{i,j+1},\ z_{i,j-1},\ z_{i+1,j},\ z_{i-1,j}$ are defined as its neighbors. The potential function for a this neighborhood is defined as:

$$H(Z) = \sum_{i,j} \ -z_{i,j}\ (z_{i,j+1}\ +\ z_{i,j-1}\ +\ z_{i+1,j}\ +\ z_{i-1,j}) \tag{6}$$

   We do not for a moment consider the Ising model a meaningful representation of porous media. We include it only because it is simple, widely understood, and therefore a convenient point of comparison.

2. *Histogram model* [2]: This model is non-parametric, keeping the entire joint probability distribution of local set of pixels within a neighborhood. We have chosen the neighborhood of a pixel to be the eight adjacent pixels, leading to a non-parametric model containing of a histogram of $2^9 = 512$ probabilities. For reasons of convenience we compute two histograms, one for configurations with a black pixel at the center and another for white. Fig. 4 shows such histogram, learned from the image in Fig. 1(a).

3.3 Gibbs Sampling Subject to Constraint

It is hardly possible to sample directly from the posterior distribution since the configuration space for an image with N sites has 2^{N^2} elements. When the random field is considered to be continuous, hierarchical and multi-scale approaches [7] can be applied. However, for the discrete case Monte Carlo Markov Chain (MCMC) methods [12] are most fitting.

The well known Gibbs sampler [10], based on an MCMC approach, generates random samples from the Gibbs distribution by producing a Markov chain whose elements are a sequence $Z_1,\ Z_2,\ ,...$, such that Z_{i-1} and Z_i can differ in at most one pixel [11].

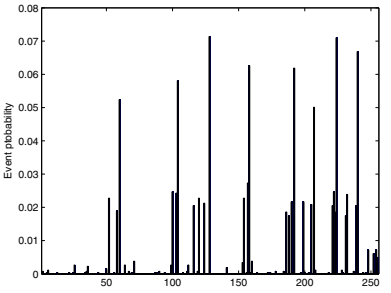

(a) Probability of 256 possible configurations which have black pixel at the center

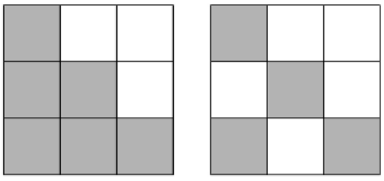

Configuration # 155
Probability = 0.024

Configuration # 78
Probability = 0.000

(b) Examples of two configurations and their probability

Fig. 4. An example of probability distribution of all possible configurations with black pixels at the center (a), with examples of two configurations and their probability (b). The random field is modeled non-parametrically via the probability of every configuration of the eight binary pixels surrounding a central pixel.

In principle, the posterior sampling appears straightforward: specify the energy function and run the Gibbs sampler. In practice the problem is not at all straightforward.

The process of annealing involves generating a sequence of samples, applying the Gibbs sampler while gradually increasing β (decreasing T) in the energy function. This annealing process is started at small β (high temperature), where $p(Z)$ is only a weak function of Z, thus Z is relatively unconstrained, and as β increases the system is driven to lower energy until the minimum energy, the most probable Z maximizing $p(Z)$, is obtained.

But here lies the problem:

1. We are not seeking the most probable realization, rather we want a *random* realization faithful to the prior model and measurements, that is, we wish to draw a random sample from $p(Z)$ at $\beta = \infty$ ($T \neq 0$).
2. However the energy function $H(Z)$, empirically derived from porous media, is really only valid at $\beta = \infty$ ($T = 0$). That is, Z is unlikely to look like a porous medium unless the empirical constraints in $H(Z)$ are rigidly asserted.

Astonishingly, almost all porous media MCMC papers ignore the above distinction and do simulated annealing to generate "random samples" from the prior model, a procedure which succeed because the annealing process fails to find the optimum Z maximizing $p(Z)$, and finds a random, near-optimum Z instead.

Therefore, we are left with three possible approaches:

1. Posterior sampling from the Gibbs distribution, problematic as described above.
2. Maximizing the Gibbs posterior distribution [11], by simulated annealing.

3. Constrained maximization of the posterior distribution [11], also by simulated annealing.

The latter two cases are distinguished by the relative choices of weighting factors α, β in(3), (4). If α is fixed and $\beta \to \infty$ (the second approach) then we have regular annealing on a fixed model. However if β is fixed, or slowly increasing such that $\beta/\alpha \to 0$ as $\alpha \to \infty$ (the third approach), then we are annealing subject to generating samples within the following constrained space [10], [11]

$$\{Z \mid \|M - f(Z)\| = 0\}. \tag{7}$$

That is, this set contains Z which matches the low-resolution measurement perfectly.

4 Results and Evaluation

To evaluate and test the proposed approach, we have applied Gibbs sampling with hard constraint in the form of low-resolution measurement to a portion of two images shown in Fig. 1. The results for those images are shown in Fig. 5, Fig. 6 and Fig. 7 three different samples are generated for the two prior models (Ising and Histogram).

If the original high-resolution image is $n \times n$ and the measurement is $m \times m$, the defined down-sampling parameter is $d = \frac{n}{m}$. For the result shown in Fig. 5, Fig. 6 and Fig. 7, $d = 8$.

The results can be evaluated in terms of the goodness of fit to the prior and measurement models. For the reconstructed image $Z_{|M}$ from true, underlying sample Z_s, we define $J(Z_{|M}|M)$ and $J(Z_{|M}|Z_T)$ to be the goodness of fit to the prior and measurement, respectively

$$J(Z_{|M}|M) = \|M - f(Z_{|M})\|^2 \tag{8}$$

$$J(Z_{|M}|Z_T) = \|Z_T - Z_{|M}\|^2. \tag{9}$$

Satisfying the prior model perfectly implies finding Z such that $H(Z) = 0$. However, a random sample found by suboptimal annealing, will be expected to have $H(Z) > 0$. The actual numerical value of $H(Z)$ is difficult to interpret, therefore we will assess all generated samples on the basis of goodness of fit to the measurement $J(Z|M)$ and the Mean-Squared Error (MSE) $J(Z|Z_T)$ to the underlying, true high-resolution ground-truth Z_T.

For the Ising model, parameter β in Eq. (3) is estimated using the method in [6]. As a one-parameter model the Ising model is able to represent only a very limited range of structures, so the reconstruction are relatively poor.

The histogram model, although with a structure nearly as local as Ising, has many more degrees of freedom (512), and is able to give a much clearer, more convincing reconstruction.

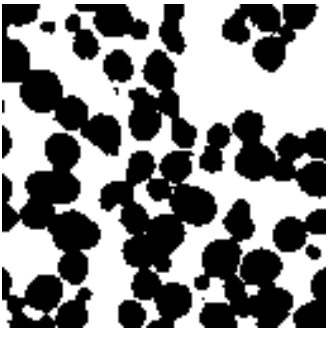

(a) Original image

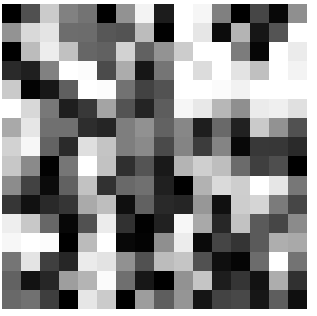

(b) Measurement

(c) Reconstructed using Ising model

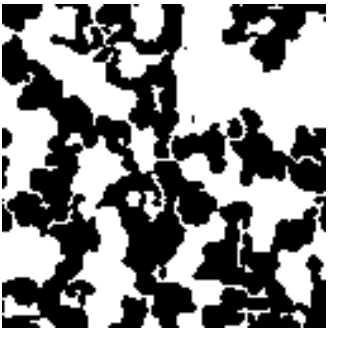

(d) Reconstructed using Histogram model

Fig. 5. reconstructed original image (a), given measurement as (b) and using two different models. The result using Ising and Histogram model is illustrated respectively in (c) and (d). The resolution of original image is 8 times greater than the measurement, i.e., d=8.

We can also study the MSE variations for different values of the down-sampling parameter d, i.e. as a function of measurement resolution. Fig. 8 shows this variation. Clearly, as expected, $J(Z|Z_T)$ is monotonically increasing with d, however the increase is relatively modest, implying that even for large d the method is producing meaningful estimates. The MSE for teh proposed method is also compared to the following reconstruction methods:

1. Random reconstruction, which assigns values $\{0, 1\}$ to the sites in the image with probability 0.5,
2. Constrained random reconstruction, which assigns values $\{0, 1\}$ to the sites in the image with probability 0.5 constrained by the measurement.

The correlation between the reconstructed and the original images is more subtle to measure. A pixel in the middle of a large pore or solid is likely to be the same in the original and reconstructed images, as opposed to a pixel on a

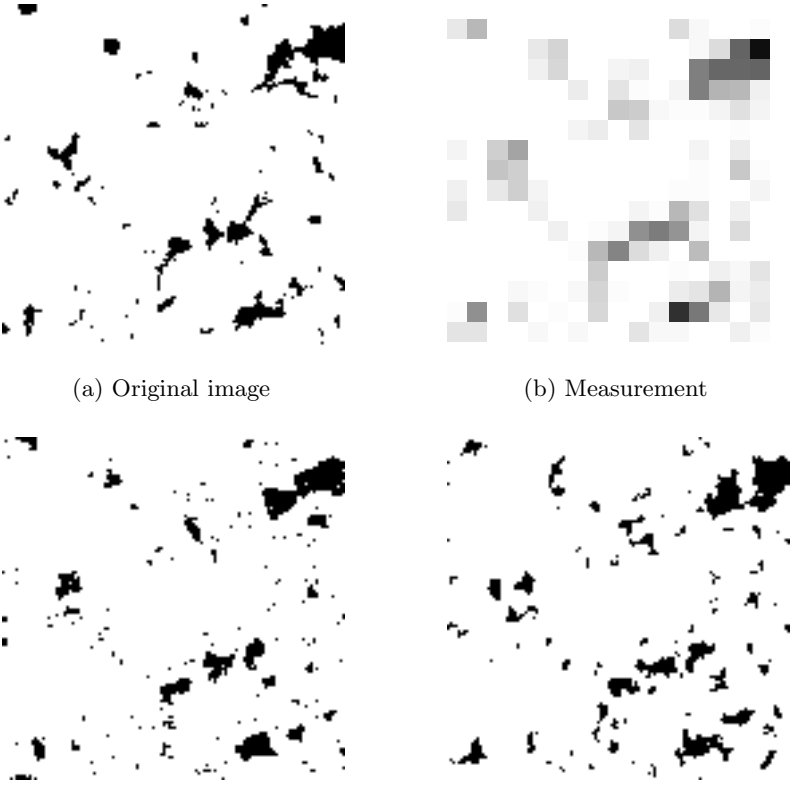

(a) Original image (b) Measurement

(c) Reconstructed using Ising model (d) Reconstructed using Histogram model

Fig. 6. reconstructed original image in (a)which is a small portion of Berea image, given measurement in (b) and using two different models. The result using Ising and Histogram model is illustrated respectively in (c) and (d). The resolution of original image is 8 times greater than the measurement, i.e. d=8.

pore/solid boundary. That is, the original-reconstructed correlation is likely a function of structure size. If we measure structure size at a pixel as the number of times a pixel value is unchanged by sampling, then we can compute correlation as a function of size. The results of Fig. 9 confirm our expectations, and quantify how many scales below the measurement resolution the posterior samples accurately reflect the measured sample.

A final comparison examines pore statistics as a function of pore size. Given an $(n \times n)$ image Z, we define $Z^{(k)}$, $k = 1\ ,\ \ldots\ ,\ n-1$, as

$$Z^{(k)}(i,j) \ = \frac{1}{k^2} \sum_{h=i}^{h+k-1} \sum_{g=j}^{g+k-1} Z(g,h). \tag{10}$$

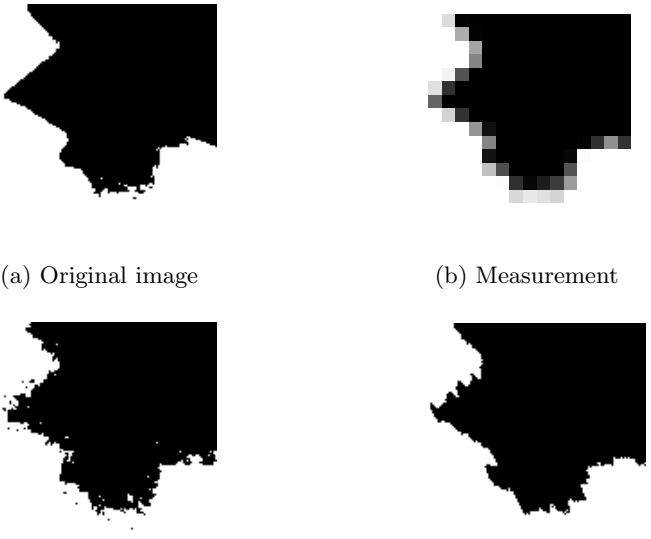

(a) Original image (b) Measurement

(c) Reconstructed using Ising model (d) Reconstructed using Histogram model

Fig. 7. reconstruction of a portion of Berea image in (a), given measurement in (b) and using two different models. The result using Ising and Histogram model is illustrated respectively in (c) and (d). The resolution of original image is 8 times greater than the measurement, i.e. d=8. Although the reconstructed image shown in (d) is more consistent with the original image shown in (c), the prior model still needs to have more contribution in the process to avoid artifacts caused by the measurement.

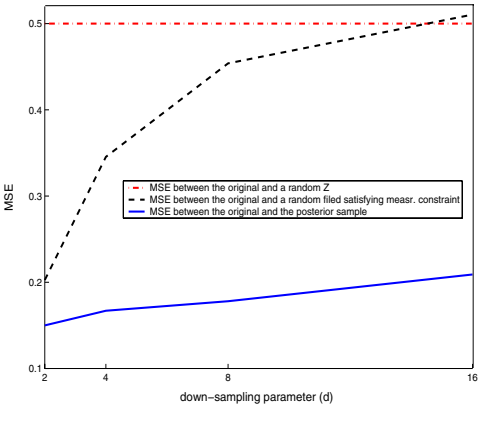

Fig. 8. How the MSE $J(Z|Z_T)$ of the reconstruction process for the image in Fig. 1(a) changes as a function of measurement resolution. $d = n/m$ is the down-sampling parameter for $m \times m$ measurement and $n \times n$ original image. The MSE $J(Z|Z_T)$ of the proposed method is also compared to two other methods: random reconstruction with probability 0.5 and constrained random reconstruction also with probability 0.5 but constrained by the measurement.

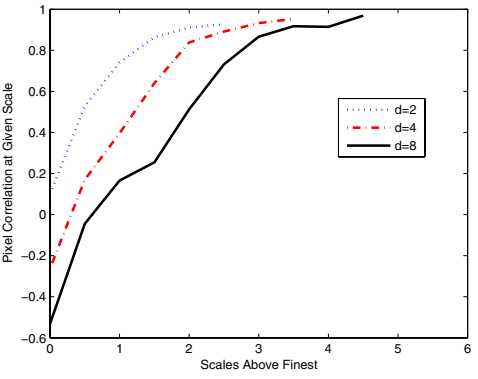

Fig. 9. Correlation between the original and reconstructed image, with the correlation computed as a function of structure scale. As would be expected, the details at the finest scales are only weakly correlated with truth, but much more strongly for larger structures.

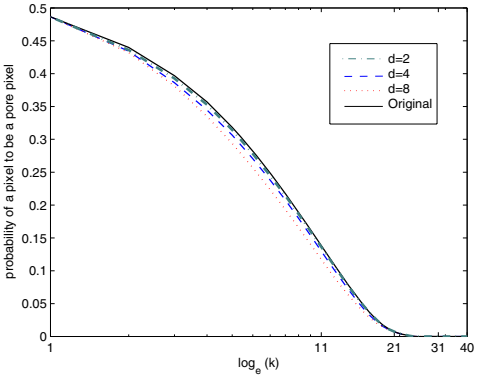

Fig. 10. Pore probability as a function of size k for the image in Fig. 1(a): By down-sampling each image into overlapping blocks of size $k = 2, 3, .., n-1$ and counting the number of pore blocks, we have a sense of pore distribution vs. scale.

In other words, each element of $Z^{(k)}$ is the sum of a $(k \times k)$ block in Z. Then we consider the probability of a pixel to be a pore pixel as the fraction of pore pixels in $Z^{(k)}$. Therefore, for each k the probability of a pixel to be a pore pixel is obtained. Fig. 10 plots the pore probability as a function of k, shown for different values of the down-sampling parameter, d. The comparison shows the effect of measurement resolution on the reconstruction process and the remarkable consistency of the reconstructed image with the pore statistics of the original.

5 Conclusion

In this paper, a model-based approach for image sampling and reconstruction of porous media is introduced. Separate from current approaches such as super-resolution and multi-resolution image fusion, we have considered a statistical model-based approach in which a given low-resolution measurement constrains the model. The approach shows considerable promise on the basis of the preservation of pore statistics and the production of meaningful structures at resolution fewer than given measurements. Comparing the Ising and histogram model in the reconstruction process shows that the later generates samples which are more consistent with the measurement and the prior model. However, a proper prior model needs to have more contribution in the process than the histogram model to avoid artifacts generated by the measurement.

References

1. P.M. Adler. *Porous Media, Geometry and Transports*. Butterworth-Heinemann series in chemical engineering. Butterworth-Heinemann, 1992.
2. S. Alexander and P. Fieguth. Hierarchical annealing for random image synthesis. *Lecture Notes in Computer Science (EMMCVPR 2003)*, (2683), 2003.
3. B. Chalmond. *Modelling and Inverse Problems in Image Analysis*, volume 155 of *Applied mathematical sciences*. Springer-Verlag, 2003.
4. Subhasis Chaudhuri. *Super Resolution Imaging*. Kluwer, 2001.
5. R. Chellappa and A. Jain. *MRF Theory and Applications*, chapter Image Modeling During 1980s: a brief overview. Academic Press, 1993.
6. H. Derin, H. Elliott, and J. Kuang. A new approach to parameter estimation for gibbs random fields. *IEEE*, 1985.
7. P. Fieguth. Hierarchical posterior sampling for images and random fileds. *IEEE, ICIP*, 2003.
8. P. Fieguth and J. Zhang. *Handbook of Video and image processing*, chapter Random Field Models. Academic Press, 2000.
9. S. B. Gelfand and S. K. Mitter. *Markov Random Fields, Theory and Applications*, chapter On Smapling Methods and Annealing Algorithms, pages 499–515. Academic Press, 1993.
10. D. Geman. *Random fields and inverse problmes in imaging*, volume 1427 of *Lecture Notes in Mathematics*. Springer-Verlag, 1990.
11. S. Geman and D. Geman. Stochastic relaxation, gibbs distribution, and teh bayesian restoration of images. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 6(6), 1984.
12. G. Winkler. *Image analysis, Random Fileds, and Markov Chain Monte Calro Methods*. Springer-Verlag, second edition, 2003.