Frozen-State Hierarchical Annealing

Wesley R. Campaigne and Paul W. Fieguth, *Senior Member, IEEE*

***Abstract*—There is significant interest in the synthesis of discrete-state random fields, particularly those possessing structure over a wide range of scales. However, given a model on some finest, pixellated scale, it is computationally very difficult to synthesize both large and small-scale structures, motivating research into hierarchical methods.**

In this paper we propose a frozen-state approach to hierarchical modeling, in which simulated annealing is performed on each scale, constrained by the state estimates at the parent scale. The approach leads to significant advantages in both modelling flexibility and computational complexity. In particular, a complex structure can be realized with very simple, local, scale-dependent models, and by constraining the domain to be annealed at finer scales to only the uncertain portions of coarser scales, the approach leads to huge improvements in computational complexity. Results are shown for a synthesis problem in porous media.

***Index Terms*—Simulated annealing, random sampling, image synthesis, hierarchical algorithms**

I. Introduction

THE synthesis of large, binary random fields has become an area of substantial interest, particularly so in the study of porous media [1]–[3], materials characterized by complex, multiscale, binary structures, for which two examples are shown in Figure 1. Although these images may appear simple, due to their binary nature, porous media and other related natural images (such as label fields in remote sensing) may in fact be near-fractal in nature, meaning that their challenge lies in the structures occupying a huge range of scales. Thus rather than continuous-state textured images (Brodatz etc.), which are more common in the literature, the focus of this paper is explicitly on the large-scale discrete-state case.

The essential challenge is how to construct a model for a given field in order to artificially synthesize further random samples, for two reasons:

1) Although it would be preferable to study *physical* samples, there is considerable time and expense associated with sample preparation and high resolution imaging. Furthermore aspects of sample preparation (cutting, polishing, exposure to air) may alter the original sample. Worst of all, the ultimate goal is to study three dimensional samples, for which thousands of repeated cutting/polishing passes is completely impractical, and three-dimensional imaging by MRI or tomographic methods so far yields samples at only very low resolution.
2) In order to analyze the macroscopic, aggregate behaviour of a material, we need multiple, *large* samples to study.

The authors are with the Department of Systems Design Engineering, University of Waterloo, Waterloo, Ontario, Canada.

We acknowledge the support of the Natural Science and Engineering Research Council of Canada.

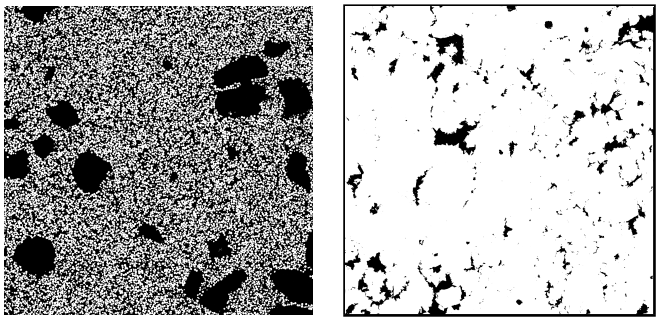

(a) Sintered glass beads (b) Carbonate rock

Fig. 1. Two examples of large 8192×8192 microscopic images having complex, multiscale structure.

Algorithm 1 Basic annealing

$k \Leftarrow 0$
Randomly initialize x_0
while not converged **do**
  $x_{k+1} \Leftarrow$ Metropolis Sampler(x_k, E, T_k)
  $k \Leftarrow k + 1$
end while
$q = k$

A great number of approaches and methods have been proposed [2], [4], [5], however all of these methods suffer from either limits on modeling complexity, such that subtle features of the porous medium cannot be represented and synthesized by the selected model, or a limit on computational complexity, such that computational complexity limits the size or accuracy of the produced sample. In this paper we report on an approach which offers tremendous improvements in both modeling and computational complexity, allowing rapid synthesis of huge multiscale two-dimensional samples, and offering the possibility of large-scale three-dimensional synthesis.

The standard approach to image sampling is based on simulated annealing [6], which is summarized in Algorithm 1, in which we repeatedly visit the pixels of a random field, using the Metropolis or Gibbs [7] sampler to update each pixel on the basis of an energy function E and an annealing temperature T_k.

In most cases, a given pixel in a random field most strongly interacts with its local neighbours, and therefore it is exceptionally difficult (i.e., slow) to synthesize structures large in size relative to the local neighbourhood.

In response to this observation, we are motivated to model a binary field in some sort of hierarchical representation, for reasons of computational efficiency, to be able to synthesize large structures in an easier fashion at a coarser scale. There are two fundamental ways in which to construct a hierarchy:

1) **Top-down,** in which the hierarchy beings at a coarse scale, with the coarse-scale elements repeatedly refined at finer scales.
2) **Bottom-up,** in which the hierarchy begins at the finest, pixellated scale, and where some sort of grouping or aggregation leads to coarser representations.

For discrete-state Markov Chain Monte Carlo (MCMC) problems [8], the bottom-up approach is considerably more common (such as Swendsen-Wang [9] and region-grouping [10] methods), in contrast to continuous-state image processing problems, in which top-down approaches dominate (wavelets [11]–[14], Laplacian pyramids [15] etc.). There are two reasons for this distinction:

1) In the continuous-state case, a coarse scale can represent a smooth, low-resolution image, which is then nudged and refined towards finer scales. However a discrete hierarchy does not allow for a smoothly-varying representation or for small refinements in state value from scale to scale.
2) In a top-down representation, a coarse-scale state element represents some square subset of the finest-scale domain. Since the finest scale will normally not be made up of piecewise-constant squares, the imposition of a regular grid from a top-down structure is not a natural fit, as opposed to the adaptive, irregular regions produced by bottom-up approaches.

Despite the above limitations and liabilities, this paper proposes a top-down hierarchical approach for the modeling and synthesis of binary random fields. Indeed, our proposed approach is able to synthesize huge 8192×8192 images possessing multi-scale structures on regular computers in a few hours of computation time.

We need to be clear that existing methods such as wavelet image synthesis, random fields synthesis using Fast Fourier transforms, and fast texture rendering methods from the computer graphics literature are all effective in their fields, efficiently producing rendered images satisfying aesthetic requirements, but which cannot however be argued to quantitatively satisfy a scientific discrete-state model. In scientific image synthesis we require a verifiable model, therefore heuristic image synthesis and enhancement methods are inappropriate.

There are two key contributions of this paper, one for each of modeling and computational complexity. First, a top-down hierarchy gives us a regular grid on a sequence of scales; the random field on each scale can be modeled by a scale-dependent model. Since even large-scale phenomena are local on a sufficiently coarse scale, it is possible to use relatively simple, *local* models on *all* scales to represent complex behaviour. Next, we propose a "frozen-state" approach, by which confident portions of coarser scales are frozen in place and cannot be modified at finer scales. This has advantages both in modeling, preventing finer-scale models from undoing or eroding structures put into place at coarser scales, and also a great advantage computationally, in that at any given scale only those "unfrozen" state elements need to be simulated.

The frozen-state hierarchy, introduced above, will be discussed in Section II. The simple, scale-dependent local model will be discussed in Section III. Modeling and computational considerations will be discussed in Sections V and VI, respectively. Finally a set of results will be presented in Section VII.

Algorithm 2 Basic Top-Down Annealing

```
Randomly initialize x_0^M
s ⇐ M
while s ≥ M do
  Iteratively anneal x_1^s, x_2^s, ..., x_q^s at temperatures T_1^s, T_2^s, ...
  if s > 1 then
    x_0^{s-1} ⇐ ρ^{s-1}(x_q^s)
    s ⇐ s − 1
  end if
end while
```

II. Frozen-state Ternary Hierarchies

For reasons of multi-scale modeling and computational efficiency, we wish to work on a discrete-state, top-down hierarchy. That is, if

$$x_{i,j} \in \{0,1\},\ (i,j) \in L \tag{1}$$

is some given, two-dimensional binary random field on lattice L, then we construct

$$x^s_{i,j},\ (i,j) \in L^s \ \text{ such that } \ L^1 \equiv L, x^1 \equiv x \tag{2}$$

that is, where scale $s = 1$ is equivalent to the finest scale. At each scale we will suppose the existence of some energy function E^s, implying a Gibbs prior [7] on x^s:

$$p\left(x^s\right) = \frac{1}{Z} e^{-E^s\left(x^s\right)} \tag{3}$$

If we define an interpolation operator

$$x^s \longleftarrow \rho^s\left(x^{s+1}\right), \tag{4}$$

then a top-down simulation proceeds as shown in Algorithm 2, whereby we begin with a coarse grid, having relatively few elements and proceed, scale by scale, to finer grids. Related approaches have been proposed in the past [16]–[18] but suffer from two particular problems:

1) From a computational perspective, the algorithm ends with some number of MCMC passes on x^0, the finest scale. Because the number of elements in x^0 is the same as in x, the computational complexity of the entire hierarchical algorithm is likely to be dominated by the finest scale, and the degree of improvement over a regular, non-hierarchical method will depend on the number of iterations needed for convergence.
2) Secondly, from a modeling perspective, the more troubling question is how to prevent the MCMC model on some scale from undoing the structures created by a model on some coarser scale. Such a structure undoing can occur for at least two reasons:
   - **Annealing Schedule:** if the MCMC annealing schedule begins at too warm a temperature, then the sampling at some scale s can, in one pass,

Algorithm 3 Ternary, Frozen-State Top-Down Annealing

Initialize $x_0^{M+1} \Leftarrow$ grey
$q \Leftarrow 0$
$s \Leftarrow M$
while $s > 0$ **do**
  $x_0^s \Leftarrow \rho^s(x_q^{s+1})$
  $G^s \Leftarrow$ grey subset of x_0^s
  $x_1^s(G^s) \Leftarrow$ random initialization
  Iteratively anneal $x_2^s(G^s), x_3^s(G^s), \ldots, x_q^s(G^s)$ at temperatures $T_1^s, T_2^s, \ldots$
  $s \Leftarrow s - 1$
end while

randomize the entire domain and undo all coarser-scale results. On the other hand, if the MCMC annealing schedule begins at too cold a temperature, the model at the current scale is prevented from making any changes, leading to mis-convergence.

- **Model Locality:** the purpose of the model E^s at each scale s is to model those phenomena *local to that scale*. Thus the model at one scale may be quite different from that at another. Therefore if the model at scale s is permitted to fully converge, the converged result is likely to be quite different from the converged result at another scale. Therefore not only the annealing schedule, but the also permitted number of iterations, must be tuned at each scale.

On the basis of the above criticisms, we introduce a novel frozen, ternary-state hierarchy, which highly effectively deals with both of the computational and modeling issues above. As opposed to the binary state of (1), we introduce an intermediate "grey" value

$$x_{i,j}^s \in \{0, g, 1\},\ (i,j) \in L^s. \tag{5}$$

The scale-to-scale interpolation operator ρ is unchanged from before, however the top-down simulation, shown in Algorithm 3, differs crucially from the previous approach in Algorithm 2, in that *only those pixels whose parents at scale $s+1$ are grey* may be changed or updated by the MCMC sampling at scale s, whereas children of parent elements which are "0" or "1" remain unchanged (frozen). Proposing such a ternary model has a huge impact on computational complexity and modeling:

- *Computationally,* now only a (possibly small) fraction of state elements at a given scale are involved in the MCMC sampling process. In many cases, the fraction of pixels which need sampling at the finest scale is a tiny fraction of the total.
- For *modeling,* because coarser "white" and "black" structures are forcibly preserved, therefore there is no need to tune the annealing schedule or number of iterations, *greatly* simplifying the annealing step at each scale, relative to previous approaches.

An illustration of the ternary hierarchy is shown in Figure 2. A given binary, finest-scale image is repeatedly sub-sampled to construct the hierarchy. The subsampling needs to be the inverse of the interpolation operator, preserving black and white regions only where the entire subset is black or white, respectively:

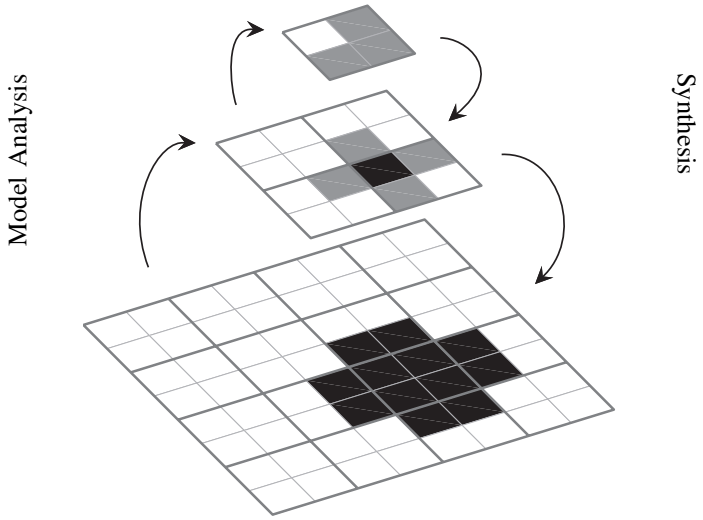

Fig. 2. The principle of ternary decomposition: a given binary image can be represented at coarser scales, such that each element on a coarser scale is selected on the basis of whether its descendants are *all* black, all white, or mixed.

$$x_{i,j}^{s+1} = \begin{cases} 0 & \text{If all children of } x_{i,j}^{s+1} \text{ at scale } s \text{ are "0"} \\ 1 & \text{If all children of } x_{i,j}^{s+1} \text{ at scale } s \text{ are "1"} \\ g & \text{Otherwise} \end{cases} \tag{6}$$

A more comprehensive illustration is shown in Figure 3, which shows the coarsification of two images, one containing relatively large spheres, and the other small spheres. We immediately observe some striking behaviours with corresponding observations:

1) The representation becomes increasingly "grey" at coarser scales, as the size of the coarse scale pixels begins to exceed the size of finest-scale structures.
   $\Rightarrow$ *Therefore the model is easily initialized as all-grey at some sufficiently coarse scale.*
2) The appearance of white seed pixels, forming the centres of finest scale spheres, takes place on a scale dependant on the sphere size. Thus the large spheres appear on scales 5 to 6, whereas small spheres appear on scale 4.
   $\Rightarrow$ *The model is inherently scale sensitive, and is naturally able to model scale-dependent structures, as evidenced by the very different statistical patterns in the two rows in Figure 3.*
3) At fine scales, the grey pixels lie only on the interface between white and black.
   $\Rightarrow$ *In most cases the interface between white and black will be only a tiny fraction of the total number of pixels. Therefore only a small fraction of the total number of pixels need to be simulated at the fine scales, precisely those scales having large numbers of pixels.*

With the ternary hierarchy defined, in principle, the following sections describe the proposed histogram model, followed

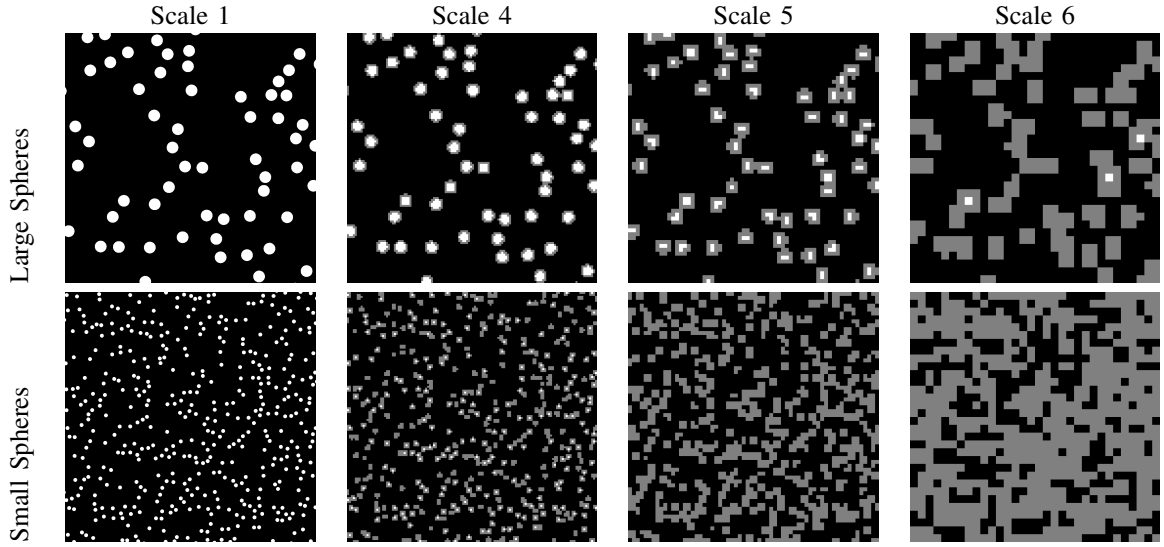

Fig. 3. Two illustrations of downsampling a given binary image per the ternary representation of Fig. 2. At coarser scales the downsampled image will become increasingly "grey", as the size of individual coarse pixels exceed the characteristic lengths of structures at the finest scale.

by results outlining both the modeling and computational advantages of this method.

III. Histogram Model

We have established the principle of a ternary hierarchy, such that at any given scale white and black structures are frozen or fixed in place, with only a grey subset remaining to be determined. These grey pixels may become black or white, and therefore frozen at finer scales, or may remain grey.

What is required is a statistical model which is compatible with describing the behaviour of an subset of grey pixels in a hierarchy. The key insight is that the hierarchical structure allows a great simplification to the model: because coarser scales have modeled and created any structures which were apparent at those larger scales, at some intermediate scale the statistical model is responsible only for *relatively local* phenomena. That is, *non-local* phenomena, which are large relative to the current scale, will have been modeled in a *local* way on some coarser scale.

Therefore rather than non-local models, based on chordlengths or correlations [2], [4], we are free to choose a comparatively simple model dependent on only local patterns. Ignoring Ising/Potts models, which are unable to represent even local patterns or textures, the simplest possible approach, known as a local binary pattern [19] (a generalization of co-occurrence matrices [20]), is essentially to preserve the joint histogram of a given pixel and its neighbours.

That is, is we define the neighbourhood of $x^s_{i,j}$ to be

$$\mathcal{N}^s_{i,j} \subset L^s \tag{7}$$

then our model needs to store the conditional probabilities

$$\Pr(x^s_{i,j} | x^s_{k,l}, \forall (k,l) \in \mathcal{N}^s_{i,j}) \tag{8}$$

The model we have adopted is a modification of the local histogram model discussed in [16], [18]. Based on our tests, it is adequate to define the neighbourhood to contain the nearest eight pixels, from which it follows that the model at each scale s is represented by $3^9 = 19683$ probabilities, which can be plotted as a histogram. An exception takes place at the finest scale, which is forced to be binary, in which case $2^9 = 512$ probabilities are stored.

An illustration of such scale-dependent histograms can be seen in Figure 4. The horizontal axis is somewhat arbitrary as it contains, in no significant order, the 19683 possible permutations of nine ternary pixels. After a brief study, however, a very clear structure emerges from the seemingly irregular histograms. At the finest scale (top), the vast majority of nine pixels are either all black (left) or all white (right). At the fourth scale, most of the domain is still all-black, however significant probabilities can be seen associated with individual white (seed) pixels. As we proceed to coarser scales the white pixels disappear entirely, the likelihood of all-black regions decreases, and grey becomes more prevalent.

Let $\bar{H}^s$ represent the learned histogram model at scale s, and $H(x)$ an operator returning the histogram of x. We wish to develop an energy function, for use in annealing, which penalizes the difference or inconsistency between a simulated field x^s and its target histogram:

$$E^s(x^s) = \|H(x^s) - \bar{H}^s\| = \sum_i \frac{(H_i(x^s) - \bar{H}^s_i)^2}{\sqrt{\bar{H}_i} + \epsilon} \tag{9}$$

where the denominator serves to allow for greater flexibility in large histogram entries, and ϵ is a small term to prevent division by zero for entries of zero in $\bar{H}$.

A slight modification to (9) is needed, as illustrated in Figure 5. Specifically, those portions of x^s that are frozen, structures created at scales coarser than s, should not be involved in the considerations at scale s, and therefore those portions of the field should not be included in the calculated histogram.

Therefore, if

$$G^{s+1} = \{i | x^{s+1}_i = g\} \tag{10}$$

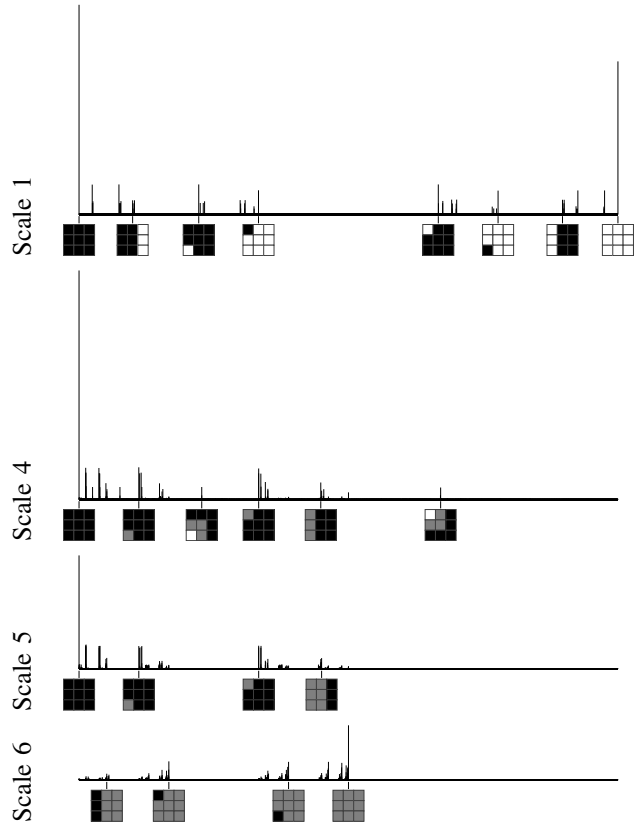

Fig. 4. The histogram model data for selected scales from the small spheres training image in Figure 3.

represents the grey elements of the parent field x^{s+1}, then

$$G^s_\rho = \rho^s(G^{s+1}) \tag{11}$$

represents those elements of x^s which are un-frozen, which are eligible for annealing. However because a given pixel contributes to the histogram entries of its eight neighbours, we need to consider the elements of x^s in $\hat{G}^s_\rho$, those elements of G^s_ρ plus their neighbours, as illustrated in the right panel of Figure 5. Therefore both the target and observed histograms are computed over $\hat{G}^s_\rho$, and (9) is modified as

$$E^s(x^s) = \left\| H\left(x^s(\hat{G}^s_\rho)\right) - \bar{H}^s \right\| \tag{12}$$

At this point we have an energy function $E^s(x^s)$ as a function of scale s, based on the scale-dependent histogram learned from training images, as in Figure 3.

IV. Methodology

The overall use of the proposed method, in practice, is to build a model from some training image, y, by iteratively coarsening it according to (6) until we reach a scale where y^s is entirely undefined ("grey"), since beyond this scale there is nothing further to model. We define this coarsest scale as $M+1$. For each scale, $s \in 1, 2, \ldots, M$, we need to acquire and store a target histogram for that scale, $\bar{H}^s$, using

$$\bar{H}^s = H\left(y^s(\hat{G}^s_\rho)\right). \tag{13}$$

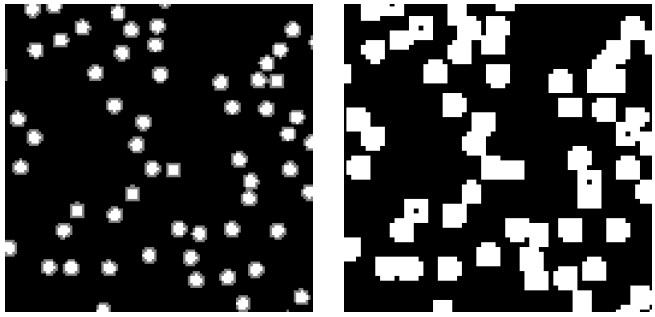

(a) Mid-Scale Image x^s (b) Counted Pixels $\hat{G}^s_\rho$

Fig. 5. In annealing down the hierarchy of scales, at a given scale only the grey pixels in (a) may be changed; all white or black pixels are frozen. Therefore only those pixels that lie inside the histogram neighbourhood of a grey pixel, i.e., the pixels highlighted in (b), enter into the energy function at that scale.

The histogram set $\{\bar{H}^1, \bar{H}^2, \ldots, \bar{H}^M\}$ forms the model input into the random sampler. Using these target histograms in conjunction with (12) as the energy function, we can then run Algorithm 3 to produce random samples through simulated annealing.

The formal convergence properties of the histogram energy function are unknown. While it has been proven that a logarithmic cooling schedule would, in principle, guarantee a global minimum [7], such a schedule is neither practical nor strictly desired. For all of the results presented in this paper, an exponential cooling schedule was used with decay factors ranging between 0.99975 at the coarsest scale and 0.992 at the finest. At each scale, the annealing was started at a super-critical temperature and continued until full convergence, which we define very conservatively as 25 iterations with no state changes.

V. Advantages in Modeling

Because large-scale and small-scale structures emerge at different scales, and with a *different* model learned at each scale, we do *not* require a single model to capture all phenomena simultaneously.

Indeed, because the ternary hierarchy freezes large-scale structures determined by coarser scales, the model at a given scale is responsible *only* for the *new* structures at the given scale, those previously unresolved at coarser scales. Therefore the model at a scale is in no way responsible to preserve large-scale structures, as they are frozen, nor does it even need to be consistent with the coarser models. This enables the use of arbitrarily high starting temperatures are each scale. The model at each scale is iterated to full convergence. Because each scale is fully converged, the temperature schedule in no way needs to be finely tuned.

Furthermore, in the ternary representation it is only the descendants of grey pixels from the parent scale which need modeling, as all white and black regions have been committed and frozen. As suggested by the images in Figure 3, we frequently observe large regions of grey having a simple structure (coarse scales), *or* small grey regions with complex structure, but only rarely do we have complex, non-local behaviour for

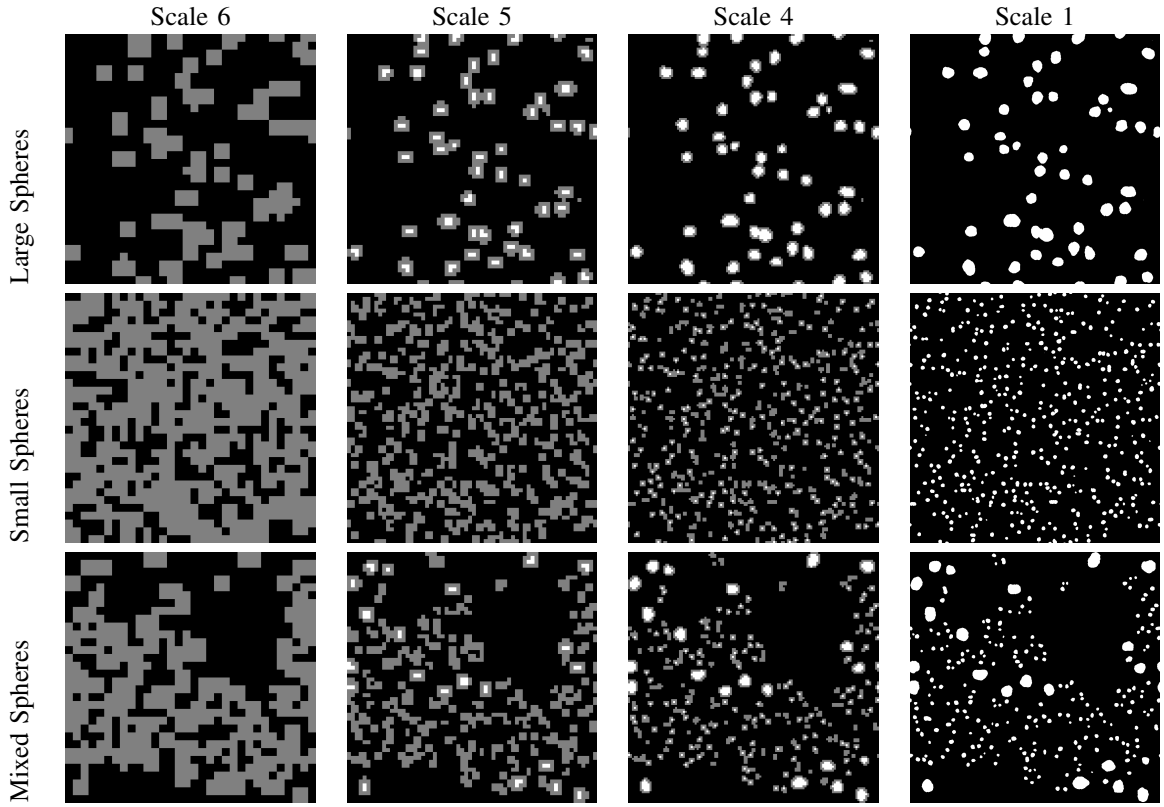

Fig. 6. Ternary hierarchical sampling results, with models learned from the large and small spheres images of Figure 3. The rather complex image of mixed spheres in the bottom row has a trivial model, nothing more than the average of the models from the top two rows. The reader can observe large spheres being seeded in scale 5 in the top and bottom rows, and the small spheres seeded in scale 4 in the bottom two rows. Such a heterogeneous, hybrid model is a huge advantage of the hierarchical representation.

grey. Therefore at every scale it is a comparatively simple, *local* model that is needed.

Such model locality in hierarchical representations has been observed in the past. In particular, a past approach to hierarchical sampling [17] asserted a chordlength model [2], [4], in which the model tests the distribution of the lengths of black and white chords. The chordlength model was truncated to test only chords up to a length of n pixels, and the reconstructed samples were found to be very nearly independent of n, for all $n > 10$ pixels. In other words, only the local portions of the model were significant and instrumental in sampling.

We maintain, therefore, that there is a highly complementary relationship between hierarchical models and simple, local energy functions. As opposed to a single-scale model, which needs to *explicitly* account for structure size by encoding the size of an object into the model, in a hierarchical model the size of an object is represented *implicitly* on the basis of the scale at which it appears.

Although such an implicit representation may not seem impressive, imagine constructing a statistical model allowing a mixture of large and small spheres. A single-scale model becomes rather complex, possibly involving hidden layers or additional labeling, whereas hierarchically the sum of large and small sphere models leads to the creation of seed pixels at both large and small scales, thus leading to the interesting, heterogeneous result shown in Figure 6. Such a hybrid model can be formed, in general, as long as the two structures being mixed are scale-resolved, since their behaviour is actually asserted by separate models on separated scales.

VI. Advantages in Computation

There are very compelling computational benefits, in many ways paralleling the preceding discussion on modeling. The frozen hierarchy specifically leads to three key computational benefits:

1) The method does not need carefully tuned annealing schedules;
2) At each scale, it is only *local* structures which need to be synthesized;
3) The number of pixels to sample at each scale is modest.

First, in terms of annealing schedules, because the large-scale structures from coarser scales are frozen in place, therefore the current scale is *unable* to destroy those structures, therefore there is no problem in starting with a hot annealing temperature. The implementation of this method does not, therefore, rely on many repeated runs to tune annealing parameters.

Second, the relatively local structures being synthesized at a scale mean that state information does not need to propagate long distances. Because an iterative local model essentially implies information being communicated via a random walk, the number of iterations required is roughly quadratic in the scale of structure to be simulated [21], so simulating only *local* structures leads to a huge reduction in iteration count, and therefore relatively quick annealing schedules can be used,

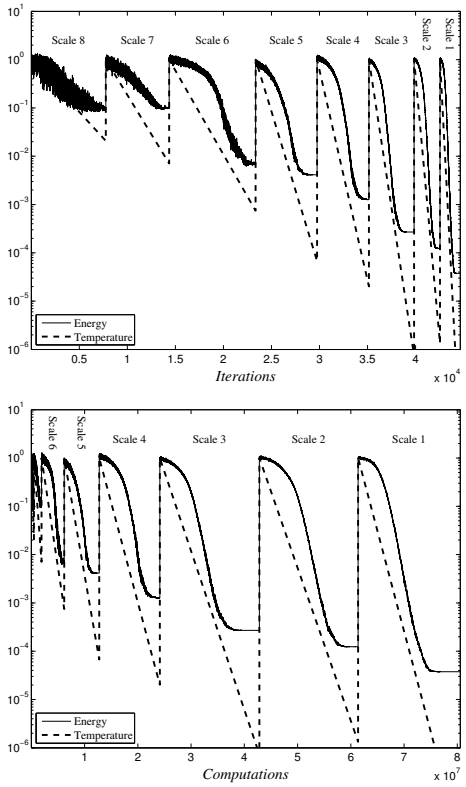

Fig. 7. The energy behaviour for the "large spheres" result of Fig. 6. The starting values of energy and temperature at each scale have been normalized. A convergence criterion (25 iterations with no changes) is used to determine when to stop annealing at each scale. As scale proceeds from coarse to fine, the computational cost of each annealing iteration across the image increases while the importance an individual pixel has on the overall result decreases. Thus, to balance computation time where it is most needed, the cooling rate is increased for finer scales.

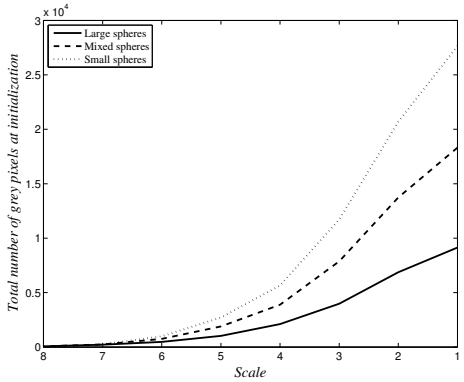

(a) Number of pixels to simulate

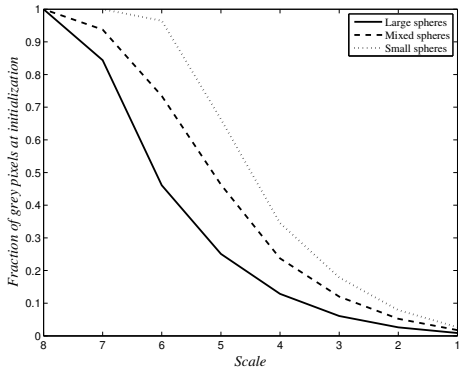

(b) Fraction of pixels to simulate

Fig. 8. The number and fraction of pixels to simulate as a function of scale, for the small/large/mixed spheres results of Fig. 6. The fraction of pixels being simulated decreases to a tiny proportion of the whole as the scale becomes finer. Although the overall fraction of white pixels is the same in all three, they are more condensed for the large spheres — fewer pixels lie on boundaries between black and white — and consequently its fraction decreases more quickly. The curves for the mixed spheres result is, unsurprisingly, consistently halfway between the other two.

as can be seen in the top panel of Fig. 7. In particular, Fig. 7 emphasizes the rather small number of iterations required at fine scales, a benefit following directly from the hierarchical nature of our modeling approach.

Next, in the ternary representation it is only the descendants of grey pixels from the parent scale which need sampling, as all white and black regions have been committed and frozen. As is implied by the images in Figs. 3 and 6, the grey pixels normally represent a declining fraction of the random field as we head to finer scales. Indeed, as can be seen in Fig. 8, for all three "spheres" test cases, the fraction of pixels to sample reduces to a tiny fraction of 1% to 5% at the finest scale. Therefore at the finest scale we may simulate only a *tiny* fraction of the pixels, a major reduction in complexity from non-hierarchical and other non-frozen standard hierarchical approaches, which still do some simulation of the entire finest scale. The bottom panel of Fig. 7 shows the modest fraction of computational complexity spent at the finest scale, despite the fact that the finest scale possesses more pixels than all other scales combined.

The overall reduction in complexity is the product of the reduction in iterations and the reduction in the fraction of pixels to be sampled. The consequence of the above improvements is enormous, relative to standard single-scale annealing. Problems which take days or weeks to converge with standard annealing, or which require carefully tuned sampling patterns and energy functions to make convergence possible at all, converge in minutes to hours with our proposed approach. A number of examples will be given in the following Results section.

VII. Results

We have simulated 64-million pixel (8192×8192) images on standard PCs in two hours, opening the possibility for far larger images, or large three-dimensional simulation on larger workstations. Figures 9 and 10 compare two such 64-million pixel samples with their training images. The similarity between the sampled and original microscopic images is stunning, preserving much of the structure on multiple scales. Comparing the two-point correlation and chordlength distributions of the samples with the originals (Figs. 9(c,d) and 10(c,d)) offers an objective measure for the quality of result, as the models do not explicitly contain either distribution. Figure 11 shows how these two samples compare with each other in the number of pixels to simulate — and, thus, the overall computational effort needed — as a function of scale. Because the sintered glass beads have a much greater pore-solid interface than the carbonate rock, the computational

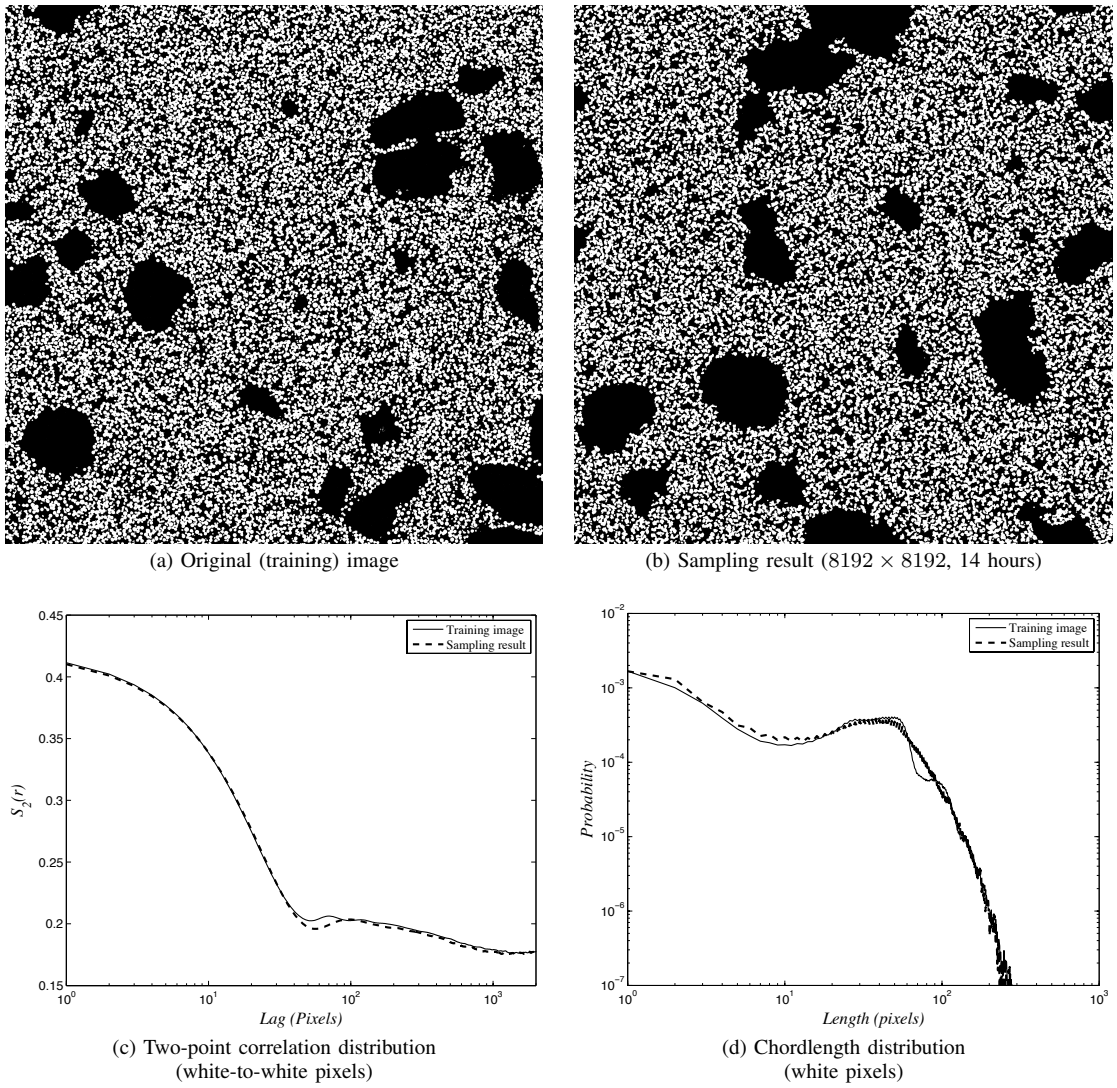

(a) Original (training) image

(b) Sampling result (8192×8192, 14 hours)

(c) Two-point correlation distribution (white-to-white pixels)

(d) Chordlength distribution (white pixels)

Fig. 9. Comparison of a full-size sampling result with the original image for sintered glass beads. The morphology of the simulated image captures behaviour over an incredibly wide range of scales, as confirmed by both the two-point correlation and chordlength distributions. The only notable deviation occurs at a chordlength of roughly 70 pixels, which corresponds to the maximum diameter of the glass beads: a hard limit which is observable only at the finest scale with a wide spatial view — precisely the type of property that a hierarchical model composed of local, scale-dependent observations is unable to capture.

burden is somewhat greater in the glass-bead case. In general, the time required to generate a sample is dominated by the number of pixels to simulate at the finest scale.

Two other examples give explicit comparisons with past methods. Figure 12 compares our proposed frozen-hierarchical approach with regular flat annealing and with a previously-proposed hierarchical approach [18]. Even after *days* of computing time, the flat annealers are, for all intents and purposes, unable to converge. Past hierarchical approaches are far more convincing in their convergence, however their sensitivity to annealing schedule means that fine scales have the ability to destroy coarse-scale structure, and therefore there is a constant challenge in parameter tuning, leading to pixellation and noise effects.

Next, to give a more comprehensive comparison to the results of [18], Figure 13 gives a zoom-in comparison of the performance of the hierarchical methods, relative to the original 8192×8192 sample. The strength of hierarchical modeling is clear here, in that the details of fine-scale boundaries require a completely different model from the large-scale morphology, and the hierarchical approach offers such scale-dependent models.

A further comparison to real textures is shown in Figure 14, based on three binarized Brodatz textures [22]. The quality of reproduction varies from excellent (D75) to poor (D108). The complexity and density of structure in D104 and D108 causes the ternary representation to become all gray in only a few levels of coarsification, meaning that only very limited large-scale structure can be represented. In particular, because our method captures only local behaviour at each scale, it has difficulty in representing highly-correlated nonlocal details, such as the long lines in D108. Where the texture fits our proposed structure, in D75 and to some degree in D104, the synthesis is effective.

Lastly, we seek to test the potential of the proposed method to be used in nonstationary multi-model settings. Rather than

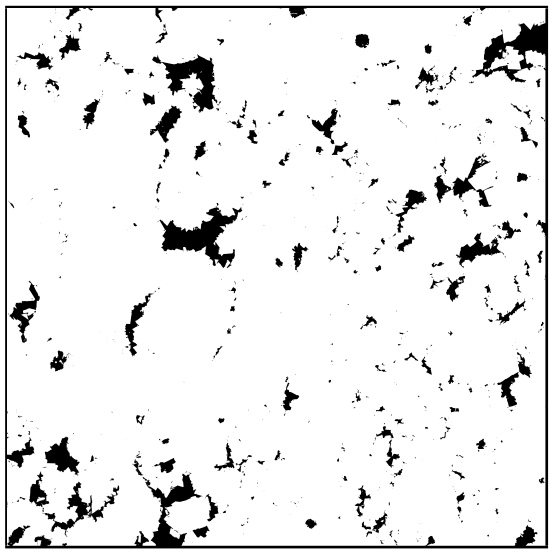

(a) Original (training) image

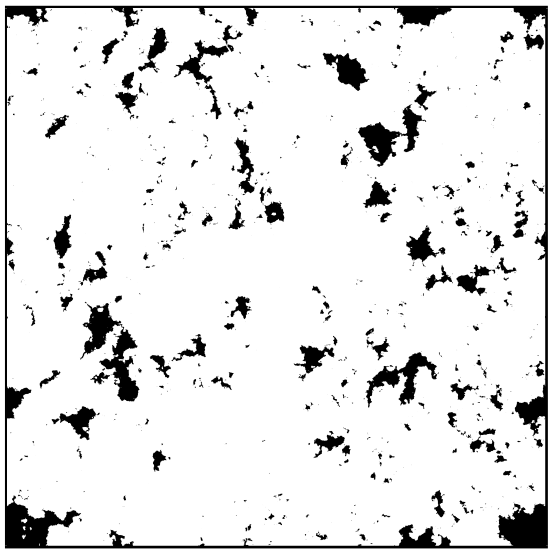

(b) Sampling result (8192×8192, 2 hours)

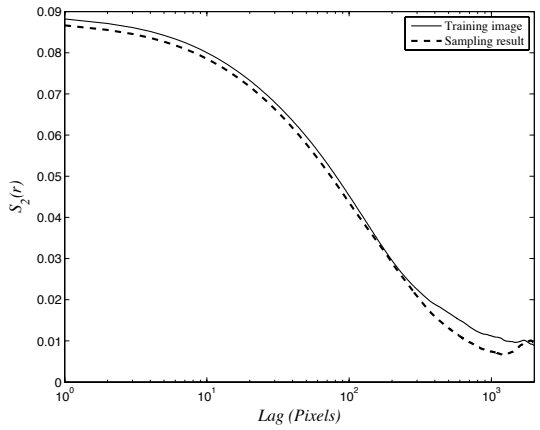

(c) Two-point correlation distribution
(black-to-black pixels)

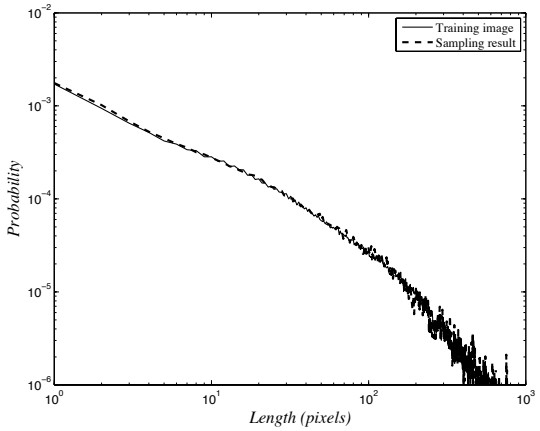

(d) Chordlength distribution
(black pixels)

Fig. 10. Comparison of a full-size sampling result with the original carbonate rock image. The sample is nearly indistinguishable from the original unless examined at full resolution. The sample and original two-point correlation and chordlength distributions match well within the expected degree of variability for other physical samples of this medium.

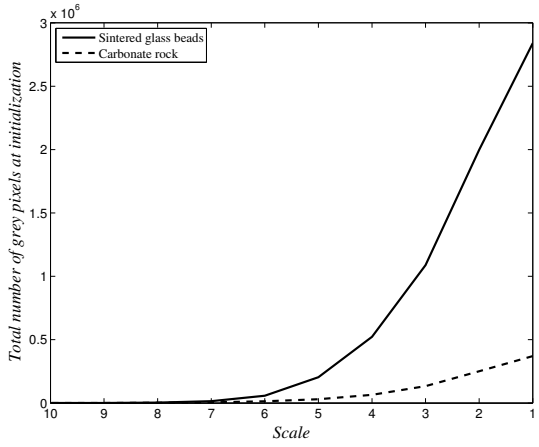

(a) Number of pixels to simulate

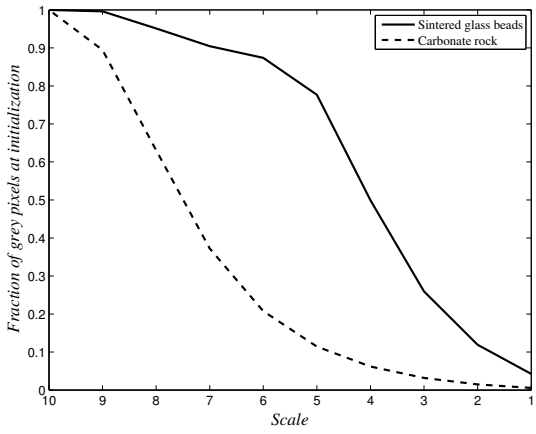

(b) Fraction of pixels to simulate

Fig. 11. Comparison of the number and fraction of pixels to simulate for the results in Figs. 9 and 10. The greater sparsity of the carbonate rock leads to a much sharper drop-off in the fraction of pixels to simulate as the scale becomes finer.

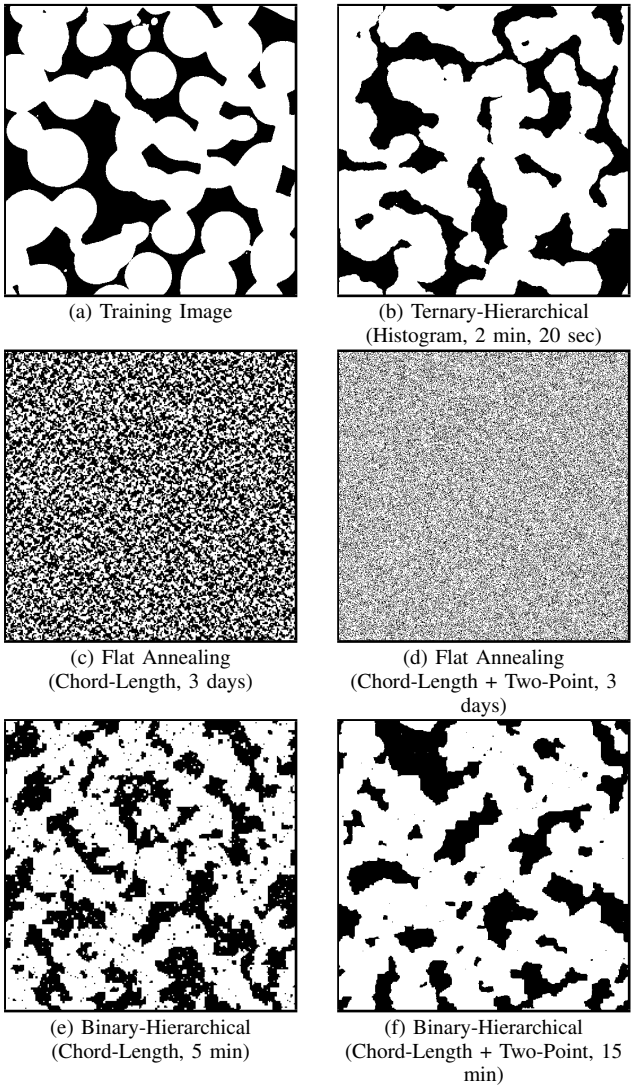

(a) Training Image
(b) Ternary-Hierarchical (Histogram, 2 min, 20 sec)
(c) Flat Annealing (Chord-Length, 3 days)
(d) Flat Annealing (Chord-Length + Two-Point, 3 days)
(e) Binary-Hierarchical (Chord-Length, 5 min)
(f) Binary-Hierarchical (Chord-Length + Two-Point, 15 min)

Fig. 12. A comparison of our proposed ternary-hierarchical method (b) with regular, flat annealing (c, d) and regular binary-hierarchical annealing [18] (e, f). Our approach outperforms the others by a wide margin, both in terms of computational complexity and in terms of reproduced morphology.

averaging two models, as was done in Fig. 6, we created two complementary models, and partitioned the physical space of the image as being represented by one model or the other. The image is still sampled all at once; the sampler, when visiting a pixel, chooses which model to invoke based on the pixel's location. The only awareness each model has of the other comes from the overlap of the local neighbourhood to pixels on either side of the model boundary, giving each model a one-pixel-width window of influence on the other — a very small overlap, but one that exists at *every scale* in the hierarchy, and so, in physical terms, starts much wider but narrows as the sampler moves from coarse to fine.

Figure 15 shows the result of this test, remarkable in the absence of any artifacts along the partition boundary. The two models are of entirely opposite natures, have no prior information about each other, and the image receives no special treatment along the boundary, nevertheless the boundary has a very natural, organic appearance. The hierarchical one-pixel overlap gives the two models sufficient influence over each other to produce a compromise along the boundary that both can accept. Such versatility to produce meaningful results, both when averaging together individual models and when sampling from models separately on a regional basis, highlights the potential for extending the frozen-state method to model phenomena of even greater complexity or nonstationarity.

The substantial advances offered by the proposed method should make it of keen interest in a variety of annealing problems. Certainly for binary porous media, the huge potential of a method able to simulate 100-million pixel domains is to perform such simulation in three dimensions, which is the subject of future work.

Acknowledgements

We would like to thank Dr. Simon Alexander for providing advice and experimental results from his binary-hierarchical annealer. We would like to thank Prof. Marios Ioannidis for porous media assistance, sample images, and his enthusiastic support.

References

[1] M. Rozman and M. Utz, "Efficient reconstruction of multiphase morphologies from correlation functions," *Physical Review E*, vol. 63, 2001.
[2] S. Torquato, *Random Heterogeneous Materials: Microstructure and Macroscopic Properties*. Springer, 2001.
[3] W. Campaigne, P. Fieguth, and S. Alexander, "Frozen-state hierarchical annealing," in *International Conference on Image Analysis and Recognition*, ser. Lecture Notes in Computer Science, vol. 4141. Springer, 2006, pp. 41–52.
[4] M. Talukdar, O. Torsaeter, and M. Ioannidis, "Stochastic reconstruction of particulate media from two-dimensional images," *Journal of Colloid and Interface Science*, vol. 248, no. 2, pp. 419–428, 2002.
[5] A. Tsakiroglou and M. Ioannidis, "Dual-porosity modelling of the pore structure and transport properties of a contaminated soil," *European Journal of Soil Science*, vol. 59, pp. 744–761, 2008.
[6] S. Kirkpatrick, C. Gelatt, and M. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, pp. 671–680, 1983.
[7] S. Geman and D. Geman, "Stochastic relaxation, gibbs distributions, and the bayesian restoration of images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 6, no. 6, pp. 721–741, 1984.
[8] G. Winkler, *Image analysis, Random Fields, and Markov Chain Monte Carlo Methods*, 2nd ed. Springer, 2003.
[9] R. Swendsen and J. Wang, "Nonuniversal critical dynamics in monte carlo simulations," *Physical Review Letters*, vol. 58, no. 2, pp. 86–88, 1987.
[10] A. Tremeau and N. Borel, "A region growing and merging algorithm to color segmentation," *Pattern Recognition*, vol. 30, no. 7, pp. 1191–1203, 1997.
[11] S. Mallat, "A theory of multiresolution signal decomposition: The wavelet representation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 11, no. 7, pp. 674–693, 1989.
[12] G. Wornell and A. Oppenheim, "Estimation of fractal signals from noisy measurements using wavelets," *IEEE Trans. Signal Process.*, vol. 40, pp. 611–623, 1992.
[13] M. Yaou and W. Chang, "Fast surface interpolation using multiresolution wavelet transform," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 16, no. 7, pp. 673–688, 1994.
[14] M. Figueiredo and R. Nowak, "Wavelet-based image estimation: An empirical bayes approach using jeffreys' noninformative prior," *IEEE Trans. Image Process.*, vol. 10, no. 9, pp. 1322–1331, 2001.
[15] P. Burt and E. Adelson, "The laplacian pyramid as a compact image code," *IEEE Trans. Commun.*, vol. 31, no. 4, pp. 532–540, 1983.
[16] S. Alexander, P. Fieguth, and E. Vrscay, "Hierarchical annealing for random image synthesis," in *Fourth International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR)*, ser. Lecture Notes in Computer Science, vol. 2683. Springer, 2003, pp. 194–210.

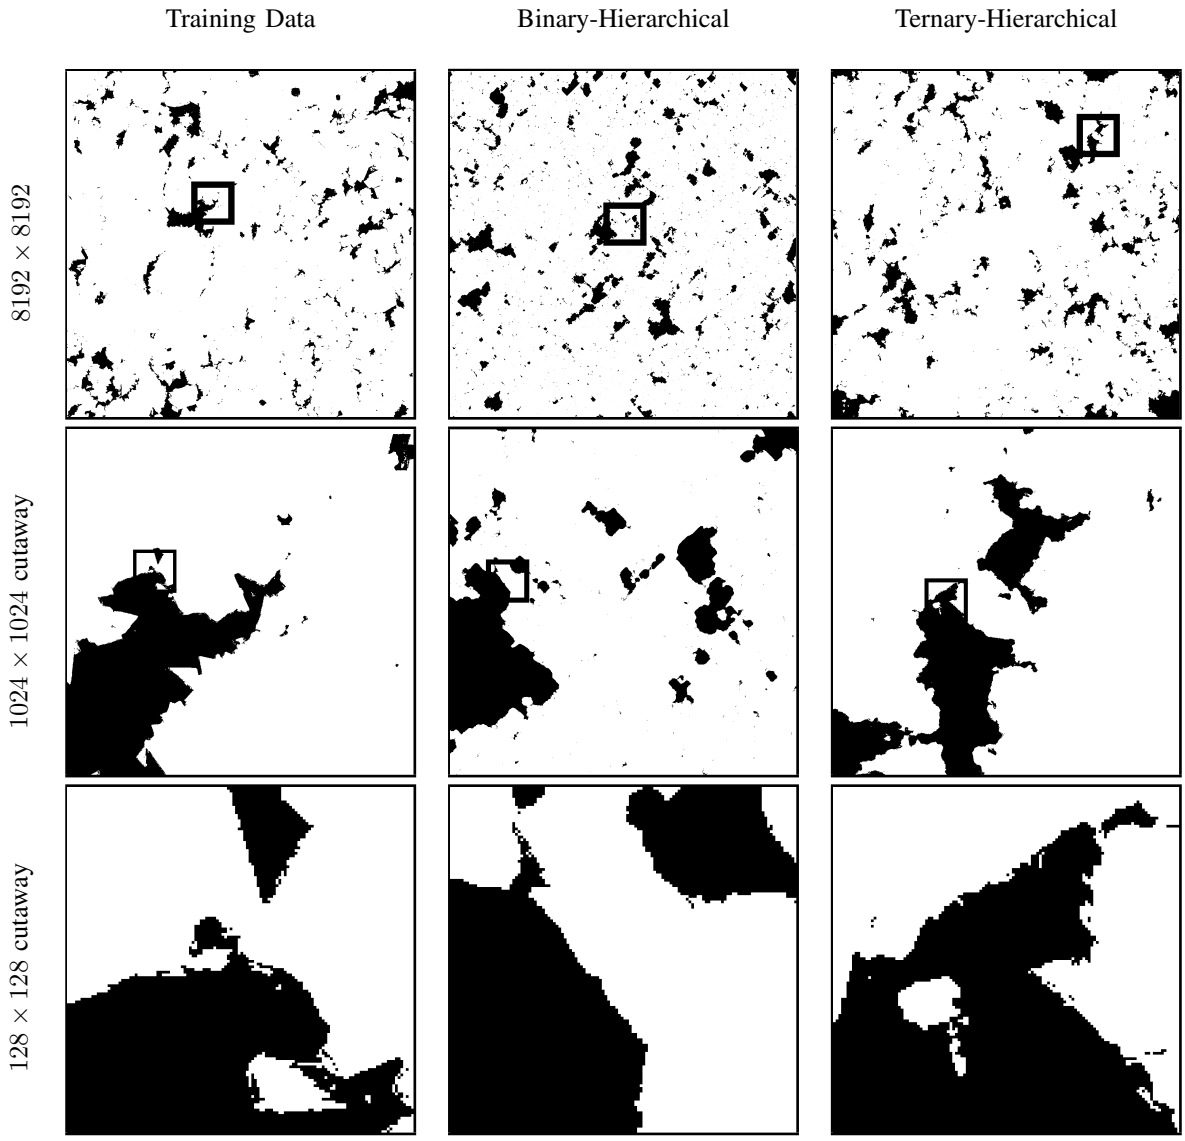

Fig. 13. A zoom-in comparison of actual data and two sets of hierarchical results for a real carbonate rock example, showing how structures are faithfully reproduced at a variety of scales. Binary-Hierarchical results from [18].

[17] ——, "Hierarchical annealing of porous media," in *SIAM Conference on Mathematical and Computational Issues in the Geosciences*, Avignon, 2005.

[18] S. Alexander, P. Fieguth, M. Ioannidis, and E. Vrscay, "Hierarchical annealing for synthesis of binary porous media images," *Mathematical Geosciences*, vol. 41, no. 4, pp. 357–378, 2009.

[19] T. Ojala, M. Pietikäinen, and D. Harwood, "A comparative study of texture measures with classification based on feature distributions," *Pattern Recognition*, vol. 29, pp. 51–59, 1996.

[20] L. Davis, S. Johns, and J. Aggarwal, "Texture analysis using generalized cooccurrence matrices," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 1, pp. 251–259, 1979.

[21] S. Wesolkowski, "Stochastic nested aggregation for images and random fields," Ph.D. dissertation, Dept. of Systems Design Engineering, University of Waterloo, 2007.

[22] P. Brodatz, *Textures: A Photographic Album for Artists and Designers*. New York: Dover Publications, 1966.

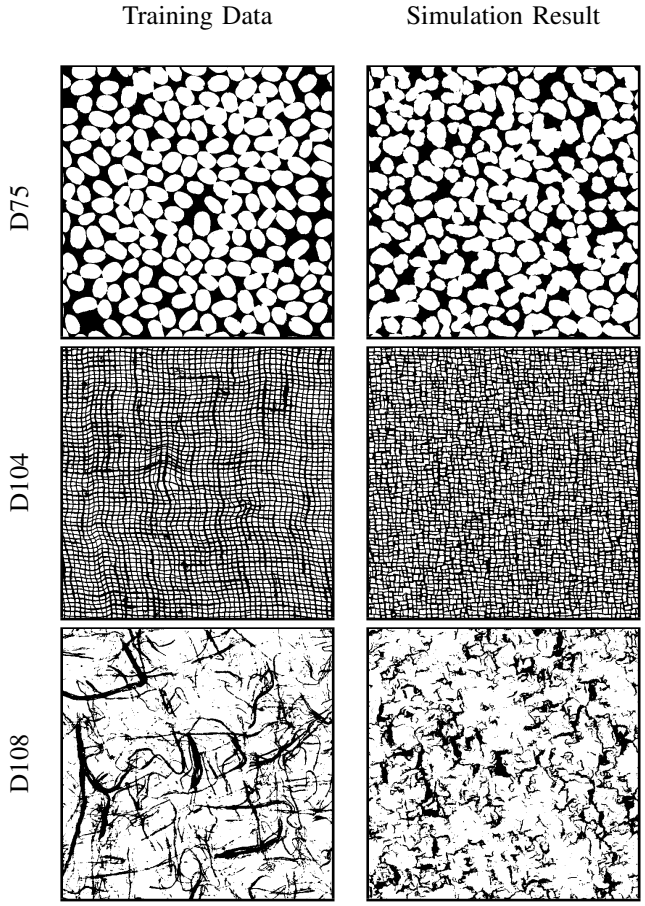

Fig. 14. An illustration of the application of our method to three binarized Brodatz textures. Since our approach models locally at each level of a hierarchy, we are unable to capture highly correlated, nonlocal, fine scale phenomena, as is clearly seen in D104 and D108, whereas the local structure in D104 and the pattern in D75 are represented well.

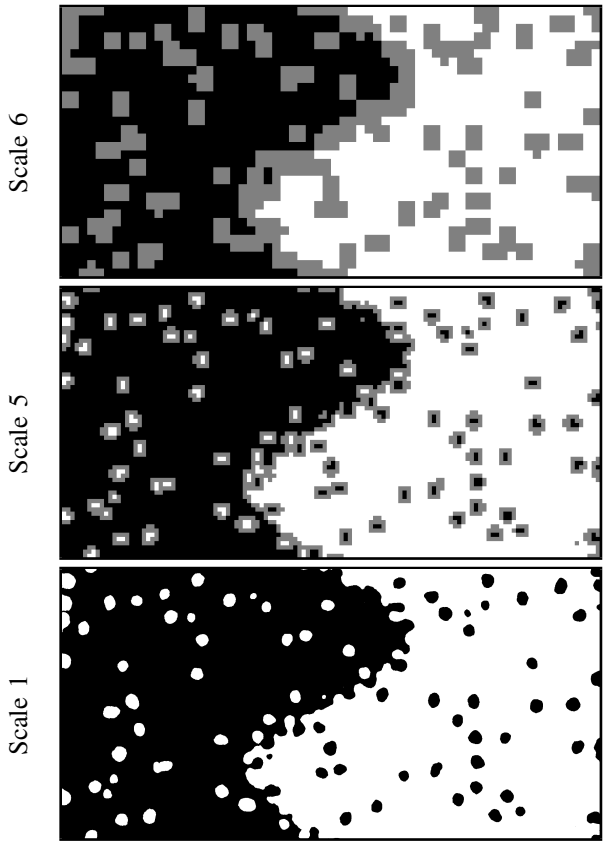

Fig. 15. Three scales are shown from a two-model example, a white-circles model on the left and its complement (black circles on white) to the right. The presence of a sinusoidal boundary separating the two model regions is evident, but discerning its exact path is difficult: both models were able to accommodate what they were able to see on the other side of the boundary, and shape their portion of it into something consistent with itself and continuous with the other side.

Wesley R. Campaigne received the B.A.Sc. and M.A.Sc. degrees in systems design engineering from the University of Waterloo, Waterloo, ON, Canada, in 2005 and 2012, respectively.

Paul W. Fieguth (S'87 – M'96 – SM'11) received the B.A.Sc. degree from the University of Waterloo, Ontario, Canada, in 1991 and the Ph.D. degree from the Massachusetts Institute of Technology, Cambridge, in 1995, both degrees in electrical engineering.

He joined the faculty at the University of Waterloo in 1996, where he is currently Professor and Department Chair in Systems Design Engineering. He has held visiting appointments at the University of Heidelberg in Germany, at INRIA/Sophia in France, at the Cambridge Research Laboratory in Boston, at Oxford University and the Rutherford Appleton Laboratory in England, and with postdoctoral positions in Computer Science at the University of Toronto and in Information and Decision Systems at MIT. His research interests include statistical signal and image processing, hierarchical algorithms, data fusion, and the interdisciplinary applications of such methods.