

Hierarchical Annealing for Synthesis of Binary Images

Simon K. Alexander · Paul Fieguth ·
Marios A. Ioannidis · Edward R. Vrscay

Received: 4 January 2008 / Accepted: 12 August 2008 / Published online: 6 February 2009
© International Association for Mathematical Geosciences 2009

Abstract This paper introduces a hierarchical annealing algorithm which addresses the very large computational costs associated with simulated annealing for the synthesis of binary images. In real-world examples, the large configuration space of such models has led to the disappointing performance of annealing approaches. Our method essentially approaches different natural scales in the image separately and in a hierarchical manner. We demonstrate orders-of-magnitude improvement compared to existing results, and discuss the inherent computational difficulties encountered with these types of approaches.

Keywords Stochastic reconstruction · Multiscale · Random field · Porous media synthesis · Random sets

1 Introduction

We report on an algorithm with very practical implications for the sampling of binary images on large grids, particularly on the construction of binary images on grids of 8192×8192 or larger, or equivalently sized 3D domains. The proposed algorithm

S.K. Alexander (✉)

Department of Mathematics, University of Houston, 651 PGH, Houston, TX, 77004-3008, USA
e-mail: simon@math.uh.edu

P. Fieguth

Department of Systems Design Engineering, University of Waterloo, Waterloo, ON, Canada

M.A. Ioannidis

Department of Chemical Engineering, University of Waterloo, Waterloo, ON, Canada

E.R. Vrscay

Department of Applied Mathematics, University of Waterloo, Waterloo, ON, Canada




Fig. 1 Examples of binary porous media images: sintered glass beads **A**, and **B** vuggy carbonate rock (high resolution, 8192×8192 pixels)

may be broadly applicable in the geosciences, where it is of interest to simulate random sets using a pixel-based approach. While we restrict demonstration of this algorithm on random sets representing the microstructure of porous materials, we note that such sets may represent hydrofacies in an aquifer, lithofacies in a reservoir, soil types in a landscape, rock types or mineralization phases in an ore deposit, or other systems of interest (Falivene et al. 2007). A great variety of porous media (Fig. 1) and composites have a random heterogeneous microstructure. The generation of 3D (or large 2D) replicas of such materials from limited (usually 2D) morphological information by stochastic methods is a problem of broad technological and scientific significance (Sahimi 1993; Torquato 2001). The significance of this problem follows from the fact that the macroscopic properties (mechanical, transport, capillary, and electromagnetic) of such materials cannot be predicted from knowledge of the volume fractions of the constituent phases alone. Instead, the geometry and topology of the convoluted 3D boundary surface separating the constituent phases must be accounted for.

Stochastic reconstruction is a valuable tool, both theoretically and practically. From a theoretical standpoint, it aids in uncovering the information content of different descriptors of the morphology of a microstructure. Such knowledge is essential for appropriately selecting the morphological descriptors to be used as reconstruction constraints (i.e., the reconstruction model) for different classes of microstructures. From a practical standpoint, stochastic reconstruction from limited morphological information can provide a domain on which to compute physical processes using fundamental equations (e.g., Navier–Stokes, heat conduction, etc.) requiring a detailed description of the geometry (solid and void phases). Much progress has been made in this direction. For example, efficient methods are now in place for computing elastic properties (Garboczi and Day 1995; Roberts and Garboczi 2002), absolute permeability (Liang et al. 2000; Spanne et al. 1994; Singh and Mohanty 2000) and formation resistivity factor (Roberts and Teubner 1995;

Ioannidis et al. 1997; Adler et al. 1992) of fluid-saturated porous media, relative permeability and electrical conductivity under two-phase flow conditions (Bekri et al. 2003), capillary pressure–saturation relationships (Magnani et al. 2000), nuclear magnetic resonance properties (Olayinka and Ioannidis 2004) and adsorption/condensation characteristics (Kainourgiakis et al. 2003) of porous solids from 3D digital representations of the microstructure.

Recently, a few general approaches to pore-scale stochastic reconstruction have been actively pursued. The first approach is based on the conditioning and truncation of Gaussian Random Fields (Adler et al. 1990; Berk 1987; Levitz 1998; Roberts and Knackstedt 1996). Though mathematically elegant and computationally efficient, this method is also model-dependent because it cannot impose constraints other than the volume fraction and the two-point correlation function. This is a serious drawback, because first- (volume fraction) and second- (two-point correlation function) order statistics are, in general, insufficient for accurately reproducing the morphology of a microstructure (Levitz 1998; Roberts 1997). Others have approached the synthesis problem using object-based approaches (Clements et al. 1990; Deutsch and Journel 1998; Gundersen and Egeland 1990; Jeulin 1997; Lehmann et al. 2008; Stoyan et al. 1996) which are computationally intractable when several orders of magnitude of length-scale are required. Finally, a general approach based on simulated annealing (SA) has been applied to the synthesis of porous media (Rozman and Utz 2001; Talukdar et al. 2002; Yeong and Torquato 1998), and to other applications in the geosciences (Deutsch and Cockerham 1994; Goovaerts 1996; Sambridge and Mosegaard 2002). This approach is model-independent, since it can impose on the reconstruction process any type and number of morphological constraints (e.g., chord-length distributions) in principle. In existing approaches, this modeling advantage is rapidly offset by the method's high demand for computational resources. For this reason, the goal of the research in this paper is to develop an alternative SA framework that preserves the modeling strengths of standard SA, but at a greatly reduced computational cost. The key to efficiency is scale.

The presence of structure at multiple length scales is a well-established fact for natural porous rocks (Radlinski et al. 2004) and the need to account for such structures in stochastic reconstructions is beginning to receive recognition (Moctezuma-Berthier et al. 2002, 2004). For example, sandstones often contain significant amounts of microporosity in addition to inter-granular porosity, whereas carbonate rocks often contain vugs and solution channels of sizes much greater than the sizes of inter-particle pores. The presence of structure at multiple length scales creates unique challenges for the reconstruction process. To cope with images of rapidly increasing resolution (Rozman and Utz 2001), the efficiency of stochastic reconstruction algorithms must increase by orders of magnitude. In an important paper, Gidas introduces the idea of renormalization to image-processing contexts, and addresses simulated annealing for the restoration of images (Gidas 1989). Our current work is naturally related, but differs primarily in two important aspects. First, Gidas discusses estimation such as the denoising of an image, in which one wishes to construct an energy surface with the ideal (denoised) image as the global minimum. By comparison, we are interested in exploring all minima of a complicated energy surface in pure synthesis. Secondly, we are not only considering proper renormalizations of our configuration

space, but also in the characteristics of a more general and hierarchical approach. This paper develops a new stochastic reconstruction algorithm that can accommodate and capitalize on the scale dependence of morphological descriptors.

In this paper, we propose a hierarchical simulated annealing (HSA) algorithm for the reconstruction of random microstructures from limited morphological information. The hierarchical scheme allows large computational gains—orders of magnitude, in cases—to be realized when applying this algorithm to models previously described in the literature (Talukdar et al. 2002). Thus, the focus of this work is on the computational benefits as applied to existing models. Although it is natural to ask about the capabilities of HSA to proposing new models, such a question is the focus of ongoing research and lies beyond the scope of this paper. The HSA approach has been successfully applied to the reconstruction (sampling) of vuggy carbonate rock at very high resolution (8192×8192) in reasonable computation times (a few minutes). This approach represents a significant and immediately practical improvement in high resolution reconstruction technique. This paper will begin with a discussion of random field models and simulated annealing (Sect. 2) to provide the needed background. Following this, our hierarchical approach is detailed in Sect. 3. The basic algorithm having been described, Sect. 4 will discuss modeling by way of energy functions. Finally, we present experimental results in Sect. 5 and some conclusions in Sect. 6.

2 Random Fields, Simulated Annealing

We propose to model our porous media images as realizations of a Markov/Gibbs random field $\mathcal{X}$, defined on a lattice in the 2D or 3D space. In order to synthesize images from this (prior) model, we wish to sample the following density

$$\pi_{\beta}(\mathcal{X}) = \frac{e^{-\beta \mathcal{E}(\mathcal{X})}}{Z_{\beta}}. \quad (1)$$

Here $\mathcal{X}$ denotes the current configuration (e.g., realization, in this case an image of a binary porous medium), where $\beta = 1/T$ is the inverse temperature parameter and $\mathcal{E}(\mathcal{X})$ is an energy function or Hamiltonian. Z_{β} is a normalization constant, known as the partition function. We are interested in drawing samples from this density, π_{β} . Due to the size of the configuration space, calculating Z_{β} directly is impractical. For this reason, a family of Markov Chain Monte Carlo (MCMC) algorithms have been developed (Geman and Geman 1984; Metropolis et al. 1953), which allow sampling from π without the knowledge or computation of Z_{β} . In what follows, we will consider the Gibbs sampler (Geman and Geman 1984). However, our approach generalizes to other sampling algorithms (Brémaud 1998; Geman and Geman 1984; Metropolis et al. 1953). The Gibbs sampler updates a single state element (image pixel) at a time, proposing a change in phase (0 to 1 or vice versa) that is accepted with some probability based on the surrounding state elements within some local neighborhood. In addition to an MCMC sampling algorithm, simulated annealing requires a cooling schedule, $\{T_k\}$, a sequence of non-negative values for the temperature parameter. As the name suggests, this may be a strictly decreasing sequence, but it is not

Algorithm 1 Simulated Annealing

```

 $k \leftarrow 0$ 
 $\mathcal{X} \leftarrow$  initial state
while  $\mathcal{E}(\mathcal{X})$  not converged do
     $\beta \leftarrow 1/T_k$ 
     $\mathcal{X} \leftarrow$  apply Gibbs sampler to  $\mathcal{X}$ 
     $k \leftarrow k + 1$ 
end while

```

necessary. While global convergence results for logarithmic cooling (Brémaud 1998; Geman and Geman 1984) exist, in practice, the slowness of the logarithmic schedule leads to other schedules being used, most commonly a geometric schedule

$$T_{k+1} = \alpha^k T_0, \quad 0 < \alpha < 1. \quad (2)$$

With the energy function, sampling method and cooling schedule defined, the standard or ‘flat’ (as opposed to hierarchical) annealing algorithm is presented in Algorithm 1.

The image synthesis algorithm uses simulated annealing to draw samples from low energy states of the distribution. However, most energy functions primarily constrain the values of proximal pixels, but many porous media exhibit structures and behavior at large scales. Intuitively, the poor computational complexity of simulated annealing approaches stems from the difficulty of building large-scale structures by the repeated, cumulative effects of local interactions. Even with relatively simple energy functions, a simulated annealing approach may be (much) too expensive. Since the computational difficulties stem from the poor fit between local interactions and large-scale structures, clearly a multiscale approach is motivated, since at coarser scales large structures become increasingly local, greatly reducing the number of iterations of local interactions. The classic example of a rigorously studied, scale-dependent binary random fields is the Ising model (Onsager 1944). However, the energy models typically encountered in porous media—much more complicated than the Ising model—preclude a similar analysis of critical slowing down or phase transition. Instead, this paper examines scale-dependent methods of decreasing the computational complexity of existing models in porous media. Figure 2 illustrates the way in which adjacent-pixel statistics can be a strong function of scale, with the coarser scales (C, F) typically less regular than the fine scales (A, D).

3 Hierarchical Annealing

We propose attacking the problem of slow convergence by using a hierarchy of scales. The construction of this hierarchy comes from two key ideas. First, most porous media contain structures at different scales, motivating a hierarchy. Second, there is a relationship between temperature and structure size. Let us expand upon the latter point. There is no absolute stability of structures during MCMC sampling;




Fig. 2 Lightly fused glass spheres (*top row*) A–C, and D–F (*bottom row*). Berea sandstone both at full resolution and two subsamplings

structures both large and small can be created and destroyed by the MCMC sampling process. However, structures of particular scales are metastable for particular temperature ranges such that, at certain temperatures, larger structures persist for a greater number of sampling iterations. This empirical observation is nothing more than the local nature of interactions. The construction or deconstruction of structures on scales beyond those of the local neighborhood can require energetically unfavorable events, which occur only with a certain probability (e.g., transition probabilities in the Markov chain). Note that, while such neighborhoods may be any clique on the fully connected graph of lattice points, we have used simple neighborhoods consisting of all points within a certain radius of the central site in practice. The larger the extent of these structures, the more such events may be needed consecutively, leading to an ever decreasing product of probabilities, roughly speaking. Consequently, large structures take a long time to build, though they may be easier to destroy—a sufficiently high temperature will simply destroy all structure in the field. The creation of this structure is clearly affected by the cooling schedule chosen; analyzing the effect of various cooling schedules has proven to be difficult (Brémaud 1998; Geman and Geman 1984; Szu and Hartley 1987).

We now consider a hierarchical approach where we want to concentrate on the synthesis of particular scales of structure at each level of the hierarchy. For an appropriate temperature range we expect the following properties relative to a particular scale: that larger scales are cold, in which the structures are metastable and slowly varying; the current scale is warm, in which the structures are variable and easily changed; finer scales are hot, in which the structures at scales finer than the current

Algorithm 2 Hierarchical Annealing

```

 $k \leftarrow 0$ 
 $\mathcal{X} \leftarrow$  initial state at coarse resolution
for scale  $s$  from coarsest to finest do
  while  $\mathcal{E}^s(\mathcal{X}_s)$  not converged do
     $\beta \leftarrow 1/T_{k,s}$ 
     $\mathcal{X}_s \leftarrow$  apply Gibbs sampler to  $\mathcal{X}_s$ 
     $k \leftarrow k + 1$ 
  end while
   $\mathcal{X}_{s-1} \leftarrow P_{s-1}(\mathcal{X}_s)$   {project to next finer resolution}
end for

```

scale are not resolved, and are essentially a random (hot) froth. These ideas taken together suggest a natural hierarchical approach. To take advantage of these observations, we construct a hierarchical approach that will concentrate, at each level of the hierarchy, only on local structure, leaving the large structures (inherited from coarser scales) intact, or mostly so, and leaving very fine details to be determined at finer scales. Since the MCMC sampler now focuses on filling in only local detail and is relieved of the burden of large-scale construction, the energy optimization at each scale is practically attainable with local interactions, not requiring large numbers of energetically unfavorable events. Algorithm 2 shows this process. In order to implement it, we need to introduce a scale parameter s to the cooling schedule (2)

$$T_{k+1,s} = \alpha_s^k T_{0,s}, \quad 0 < \alpha_s < 1. \quad (3)$$

At this point, we need to discuss the meaning of energy functions evaluated at different scales. Given an energy function $\mathcal{E}(\mathcal{X})$ on the data at its natural (finest) scale, what does it mean to evaluate the energy at coarser scales? Essentially, we have three choices. We can always operate at the highest resolution with coarse versions of the state being projected to the highest resolution where the energy is evaluated. We can infer the equivalent energy function for the coarser space from the original energy function and the method of coarsening/rescaling the space. A third alternative is to parametrize and learn the energy model at each level. The first option is impractical for computational reasons on large data sets. Solving the hierarchical problem entirely at the finest scale runs contrary to the whole motivation and spirit of the proposed hierarchical approach. The second option is normally impractical, depending on the method of projection from scale to scale, related to the idea of renormalization (Gidas 1989). If the energy function is constructed in such a way that it naturally scales with length (e.g., distribution of chord lengths) then a coarser version may readily be constructed. However, this is quite difficult to do in general, even for simple models such as 2D Ising (Schonmann and Shlosman 1995). The third method is practical since the purpose of intermediate scales is to produce the local structures appearing at that scale, it is reasonable to learn an energy function based on the empirical samples projected to that scale.

Although a proof of convergence is unavailable for such a complex process, particularly with the empirical scale-dependent energy functions for this problem domain, it should be kept in mind that essentially all practical annealing schedules in

porous media applications involve relaxation rates which are taken far faster than those for which analytic convergence properties are known. Nevertheless, our proposed method converges extremely rapidly and reliably for a variety of test cases in practice, illustrated later under experimental results. Until now, stochastic approaches like simulated annealing were understood to have real theoretical advantages in the sampling of porous media—the elegant and flexible modeling of a variety of energy functions. However, they are too simple or too expensive for practical use. Even allowing for the progress in computational hardware, simulated annealing methods are restricted to tiny grids, defeating the possibility of models capturing significant detail across many length scales. Our algorithm allows very large samples to be drawn in practical computation times. The immediate results are quite satisfying. But more importantly, this allows for ambitious modeling work in the future and the possibility of sophisticated models sampled at high resolution, even in 3D.

4 Models/Energy Functions

The purpose of this work is to describe the computational benefits made possible by using a hierarchical approach to simulated annealing. As such, any discussion of the relative merits of particular models is eschewed. Instead, we present results on common models found in the appropriate literature with little comment on efficacy or verification. In particular, there are several components of energy functions that have been used (Talukdar et al. 2002; Torquato 2001) in related approaches. We will demonstrate how the increased performance of the HSA approach improves results with these models, relative to performance for flat annealing. It gives similar or better results, with drastically reduced computational times.

In general, our energy/cost functions will be made of one or more components, possibly mixtures of multiple models, so the general finest-scale formulation will be

$$\mathcal{E}(\mathcal{X}) \equiv \mathcal{E}^0(\mathcal{X}) = \sum_i c_i \mathcal{E}_i^0(\mathcal{X}), \quad c_i > 0 \quad \forall i, \quad (4)$$

with the weighting coefficients c_i specified as part of the model. In the hierarchical case, energy functions $\mathcal{E}^s$ will be defined for every scale s , induced from the standard, finest-scale energy $\mathcal{E}^0$

$$\mathcal{E}^s(\mathcal{X}) = \sum_i c_i \mathcal{E}_i^s(\mathcal{X}), \quad c_i > 0 \quad \forall i. \quad (5)$$

Particular energy components of interest would include (Talukdar et al. 2002; Torquato 2001; Yeong and Torquato 1998): (1) the one-point correlation function, (2) the two-point correlation function, (3) the chord-length distribution, (4) the dual pore-solid chord-length distribution, (5) the lineal path distribution, (6) the “pore size distribution”, and (7) the localized distribution (histogram). It is important to recognize that all of these standard, common energy functions (none of them inherently hierarchical) can readily be implemented in our hierarchical framework, as can any

other approach that can be formulated this way. We will define the ones used to generate figures in this paper, for completeness and for readers unfamiliar with such energy functions. Denote the lattice size of the image at a particular scale s as $\mathcal{O}(s)$. We are considering binary solid/pore structures (Talukdar et al. 2002), so we denote $I^s(x)$ the indicator function for our binary image at scale s , (0 for pore, 1 for solid). Finally, let $\langle \cdot \rangle_{x \in \mathcal{X}}$ denote a spatial average (over the image).

A “One-point energy function” can be defined as

$$\mathcal{E}_1^s = \sum_{r=1}^{\mathcal{O}(s)} \|\hat{S}^s(r) - S^s(r)\|,$$

where $S^s(r) = \langle I^s(x+r) \rangle_{x \in \mathcal{X}}$ is the average density, or “one-point” correlation, and $\hat{S}^s$ is the target correlation, deduced analytically or inferred empirically from other data sets.

A “Two-point energy function” is defined as

$$\mathcal{E}_2^s = \sum_{r=1}^{\mathcal{O}(s)} \|\hat{S}^s(0, r) - S^s(0, r)\| + \|\hat{S}^s(r, 0) - S^s(r, 0)\|,$$

computed in terms of the two-point correlation $S^s(r_1, r_2) = \langle I^s(x+r_1)I^s(x+r_2) \rangle_{x \in \mathcal{X}}$.

The “Chord-Length energy function” may be approached in a similar manner. If we restrict ourselves to the horizontal and vertical directions, this is essentially the distribution of the contiguous “run length” of density pixels in these directions. Denoting these probability mass functions (estimated from the data) as p_C^h and p_C^v for the horizontal and vertical directions, respectively, we may construct an energy function

$$\mathcal{E}_3^s = \sum_{r=1}^{\mathcal{O}(s)} \|\hat{p}_C^h(r) - p_C^h(r)\| + \|\hat{p}_C^v(r) - p_C^v(r)\|.$$

The chord-length and one- and two-point energies are used in several examples in the following discussion. They may also form a component of a final energy function tuned to a particular synthesis problem (Fig. 14).

The “Local Histogram energy function” model has proven to be quite useful in practice. For a given discrete neighborhood structure, we can assign an (arbitrary) ordering on the neighborhood. Since the data is binary, we have in this way constructed a natural identification with an integers binary expansion; the n th bit state is the state in the n th neighborhood location. Figure 3 illustrates the spatial relationship for a 3rd order neighborhood in 2D, forming a map onto $0 \cdots 4095$. This forms one component of the energy function used to generate the large scale results (Fig. 14). Given this mapping, a complete state of the neighborhood local to a site is expressed as an integer in $0 \cdots 2^b - 1 = N$, where b is the cardinality of the neighborhood set. The distribution of localized neighborhoods (for a fixed neighborhood size) is then a distribution of integers. In training data, the site is either void or solid so we condition on the state of the site itself, either the probability $p_H^b(n)$ of any particular configuration

Fig. 3 Bit indices for a 3rd-order neighborhood inducing a unique mapping onto $0 \dots (2^{12} - 1)$

•	•	9	•	•
•	8	1	5	•
12	4	◦	2	10
•	7	3	6	•
•	•	11	•	•

$n \in 0 \dots N$ given the central site is black (void) or $p_H^w(n)$, and given the central pixel is white (solid). We can now estimate the distribution of neighborhood configurations in the current state from data. Like the chord-length case, an energy function can be formed by summing over all the N possible configurations

$$\mathcal{E}_4^s = \sum_{n=0}^N \|\hat{p}_H^b(n) - p_H^b(n)\| + \|\hat{p}_H^w(n) - p_H^w(n)\|.$$

This approach is discussed by Alexander et al. (2003). Note that the natural way to compare target and current distributions (approximated here by histograms) will be a Kullback–Leibler divergence. However, for practical optimization, an MSE approach is efficient and works very well. Of course, there are many other possibilities, but most of our experiments are based on combinations of these components. Note that, in principle, a modeling approach may seek to address each scale in (5) separately. Thus, our hierarchical approach admits energy functions that are not possible with a single scale algorithm. Where results are given here with energies for both hierarchical and flat annealing, all reported energies are defined on the finest scale and are directly comparable.

5 Experiments

5.1 Synthetic Data

Physical porous media samples tend to have complicated morphologies, possibly poorly modeled by simple energy functions. It can be difficult to evaluate the performance of a sampling algorithm, but for recent approaches see Arns et al. (2004). Therefore, it is informative to construct synthetic data that can be modeled exactly, to emphasize the benefits of a hierarchical approach. Preliminary versions of some of these experiments were discussed in conference talks (Alexander et al. 2004a, 2004b).

First, consider a set of 300 images generated by means of two Gaussian random fields (GRFs) (for ease of sampling), and chosen to exhibit two characteristic length scales. Figure 4 shows an example of this data set, which allows us to concentrate on two characteristic length scales and on a morphology simpler than may be found in some physical samples. Convergence is easy to see in the evaluation of the chord-length distributions. At the highest resolution, there are small isolated chords on the




Fig. 4 GRF data generated with two characteristic scales, as shown in **A**. A strong bias toward black pixels in the shorter length scale GRF has the effect that after several rescalings, the coarse resolution images such as **B** are much simpler




Fig. 5 Convergence of flat (*red*) and hierarchical (*blue*) annealing for GRF data set. The spikes in the HSA profile are due to the repeated projection from one scale to the next

edges of large features and longer chords comprising large features. In the annealing algorithm, moving between isolated small chords through isolated medium length chords to eventually collate into larger structures is not energetically favorable. We have slowing down as the structure begins to emerge, since these events are necessary, but of low probability leading to a large number of proposed steps before




Fig. 6 GRF training data. Each panel gives the final energy, $\mathcal{E}(\mathcal{X})$, and approximate runtime (on a 3 Ghz Pentium IV class machine). Shown are example training data **A**, *best case* sample from SA **B** (with very slow geometric cooling), sample from HSA method **C**, and finally **D**, sample from SA with equivalent computation to the hierarchical case. Model used was two-point probability and both white and black phase chord-length distributions. Images are 512×512 pixels

improvement. In contrast, coarse levels of the hierarchical annealing do not have any small isolated chords, so we may expect convergence to be quicker. Empirical results support this conjecture. Figure 5 compares flat and hierarchical sampling runs on a plot of energy verses computations. The energy spikes in the hierarchical annealing curve are due to projection from one scale to the next. Immediately after projection, local configurations are at a high energy due to artifacts of the projection operation. Since these high energies are due to local errors or inconsistencies, they are easily remedied and energy levels quickly drop as the sampling algorithm progresses.

Figure 6 shows 512×512 images sampled from this model using three methods. The first panel shows a Gaussian Markov Random Field (GMRF) sample, typical of




Fig. 7 Zoom views of images shown in Fig. 6

the data set we have trained the model on. This is followed by the best flat SA result allowing about 3 days computation, the best HSA result taking about 15 minutes, and the best SA with a different cooling schedule taking the same number of computations as the HSA sample. As be seen here and in a zoom view (Fig. 7), the complex morphology is not represented in the SA samples. The long-run result lacks medium-size structures, and the fast SA sample is very poor. The HSA result, while clearly less than perfect from a modeling point of view, shows both visually and by final energy a far superior result over regular SA.

5.2 Computational Benefits

Separate from issues of convergence (analysis of which is difficult for non-logarithmic cooling schedules, and more so where multiple scales are concerned), there are clear sources of expected computational benefit from the HSA approach. At




Fig. 8 Energy vs. Computation: Results for many parameterizations of the cooling schedule (hierarchical annealing)

each level of the hierarchy, we are running a Gibbs-type sampler in a smaller configuration space than at the following levels. In the cases described here, each level in the hierarchy reduces the image size by a factor of four. In 3D, the gain is much larger, of course, so we have a geometric reduction in the size of the configuration space. Beyond this simple reduction, there are data-dependent benefits. Since large scale structures are built by local interactions, the probability of constructing something when the intermediate steps are not energetically favorable is reduced as the number of steps increases (i.e., the product of the probabilities of individual steps). In a hierarchical approach, large structures are put together with fewer intermediate steps at a coarser resolution. As this is data dependent, it is difficult to quantify. However, it will benefit more from a hierarchy when the intermediate steps are less probable.

Evidence of both of these effects was already presented in Fig. 5 for our synthetic data. Figure 8 shows the results of many hierarchical runs with a simple parametrization on the cooling schedule. In these images, the data set was Berea sandstone, and the model was chord-length distributions in both phases. It is important to note that the parametrization allowed for poor choices of the cooling schedule (that would allow high enough temperatures to tear apart structure after projecting to a high resolution), and these paths are especially evident at the top-right of the ensemble. Many parameterizations allow for very good performance. Figure 9 compares multiple SA




Fig. 9 Energy vs. Computation: Several hierarchical annealing runs (*solid*) compared to six flat annealing runs (*dashed*), for a variety of cooling schedules. The flat annealing runs clearly show the trade off between computational complexity and final energy. (Increasing α trades more computations for better final energy)

and HSA results, all based on geometric cooling schedules (2). Slowing the SA cooling schedule improves the final energy, but at the cost of more computations. However, none of the SA runs are able to match HSA for both computational time and final energy. Figure 9 is presented with log–log scales and demonstrates a computation versus energy gain of more than an order of magnitude.

As previously noted, an interesting demonstration of the slow convergence of some models is an energy function made up of both positive (white) and negative (black) chord-length distributions. The following discussion and figures are not presented as a viable model reconstructing porous media. Rather, it is a very simple example that demonstrates the difficulties that may be encountered. When attempting to build large scale structures, such as in our GRF data set, there is very little energy difference made by any intermediate step. In some sense, the white and black chord distributions are at odds and therefore very slow convergence results. While this is definitely not a good model to capture the morphology in question, it is interesting that even allowing many days, flat annealing will not converge at all usefully while hierarchical-annealing generates large scale structure in a few minutes (Fig. 10). Figure 11 shows the same computational methods as Fig. 10, but with a data set of lightly fused glass beads. Note the complete failure of convergence for the flat annealing cases.




Fig. 10 Similar to Fig. 6, these panels show results for the more difficult case where the energy function is a mix of white and black phase chord-lengths

5.3 Physical Data

While empirical results on the synthetic data described in the previous section are valuable, it is of interest to apply the approach to real data. It is worth re-iterating that the purpose of this current work is not to address modeling issues in the problem domain. For this reason, validation of resultant images is difficult, at least in an absolute sense. While subtle issues of validation are beyond the scope of this paper, flat annealing improvements can be extremely clear (Figs. 10 and 11) in comparison. In contrast, Fig. 12 shows similar samples, but for a more realistic model of both chord-length distribution and two-point correlation. Although the prior model does not capture all of the morphologies of the training set, it is very clear that HSA samples are being drawn from a class of images with chord-length distributions that are strongly in agreement with the mean chord-length of the training set (Fig. 13). In general, we see that the simulated process under the HSA algorithm is converging to




Fig. 11 Sintered beads data set. Again, each panel gives the final energy, $\mathcal{E}(\mathcal{X})$, and approximate runtime. Shown are example training data **A**, sample from SA with very slow geometric cooling **B**, sample from SA with equivalent computation to the hierarchical case but with enforced volume fraction (to the mean of the training set) **C**, and finally **D**, sample from HSA method described herein. Model is (cf. Fig. 10) both white and black phase chord-length distributions. Images are 256×256 pixels

far lower energies much more rapidly than under standard SA. Figure 14 compares the training data (left) and HSA reconstruction (right) with the data shown at full resolution and with two sub-images (1024×1024 , 128×128). The sub-image regions indicated at the previous resolution. This reconstruction uses both chord-length and a local neighborhood structure approach (see Sect. 4, Models/Energy Functions). In this synthesis, energies at each scale were learned separately.

Visually, these figures demonstrate that the results look very good. Even better, the autocorrelation was not used in the reconstructing energy function and can therefore serve as a useful validation of the result. Figure 15 plots the observed autocorrelation for the two images, showing excellent local (small lag) agreement and reasonable agreement at large distances.






A sintered beads sample: $\mathcal{E} = 0.221$ B SA: $\mathcal{E} = 3.94$, 3 daysC SA: $\mathcal{E} = 11.0$, 15 minD HSA: $\mathcal{E} = 0.483$, 15 min

Fig. 12 Sintered beads data set with a slightly more realistic energy function, the mixture of two-point probability and chord-length distribution. In this case the images are 512×512

Fig. 13 Chord-Length distributions for spheres data set (shown in Fig. 12). The *dashed line* is mean chord-length distribution over training set, and *dotted lines* give one standard deviation. The *solid line* is final distribution of an HSA sample, while *dot-dash line* shows a typical sample from the training set





A 8192×8192 physical sampleB 8192×8192 HSA resultC 1024×1024 region of sampleD 1024×1024 region of HSA resultE 128×128 region of sampleF 128×128 region of HSA result

Fig. 14 High-resolution image of vuggy carbonate rock (1.87 micron/pixel) showing structure at multiple scales (*left column*) and HSA sample (*right column*). Energy function is based on local neighborhood structure and chord-length distributions. The reconstruction took approximately 4 days, with a final energy of $\mathcal{E} = 2.85e-4$ (the original image has energy of $4.44e-6$)

Fig. 15 Plots of the autocorrelation function $S_2(r)$ for both training and reconstructed images shown in Fig. 14. Although autocorrelation is not directly part of the energy function used, agreement is very good



6 Discussion and Conclusions

We have demonstrated the computational benefit of a hierarchical approach to simulated annealing in the computationally expensive approach of pure synthesis (e.g., sampling) from low temperature Gibbs distributions (1). This approach now allows researchers to attempt more difficult synthesis problems on large images with multiple natural length scales. Clearly, with such an approach, there is a modeling component not addressed in the current work. While some comments relevant to modeling in this hierarchy have been made, this aspect is the subject of another paper. Being able to model reliably and process 64-million pixel random fields obviously raises a wide variety of further questions and opens new opportunities in the synthesis of porous media images. The most immediate extension of this work is to three-dimensional data, where the large voxel count quickly leads to prohibitive computational complexity with standard SA. Furthermore, in addition to expanding the list of standard SA models implemented and tested in the HSA context, the HSA method allows for new types of scale-dependent models that is the object of ongoing research.

References

- Adler P, Jacquin C, Quiblier J (1990) Flow in simulated porous media. *Int J Multiph Flow* 16(4):691–712
- Adler P, Jacquin C, Thovert JF (1992) The formation factor of reconstructed porous media. *Water Resour Res* 28(6):1571–1576
- Alexander S, Fieguth P, Vrscaj E (2003) Hierarchical annealing for random image synthesis. In: *EMM-CVPR'03*. LNCS, vol 2683. Springer, Berlin, pp 194–210
- Alexander S, Fieguth P, Vrscaj E (2004a) Hierarchical annealing for scientific models. In: *ICASSP'04*. IEEE Press, New York, pp 33–36
- Alexander SK, Fieguth P, Vrscaj ER (2004b) Parameterized hierarchical annealing for scientific models. In: *International conference on image analysis and recognition (ICIAR'04)*. LNCS, vol 2683. Springer, Berlin, pp 236–243
- Arns CH, Knackstedt MA, Mecke KR (2004) Characterisation of irregular spatial structures by parallel sets and integral geometric measures. *Colloids Surf Physicochem Eng Asp* 24:351–372

- Bekri S, Howard J, Muller J, Adler PM (2003) Electrical resistivity index in multiphase flow through porous media. *Transp Porous Media* 51(1):41–65
- Berk NF (1987) Scattering properties of a model bicontinuous structure with a well defined length scale. *Phys Rev Lett* 58(25):2718–2721
- Brémaud P (1998) Markov Chains: Gibbs Fields, Monte Carlo simulation, and queues. Texts in applied mathematics, vol 31. Springer, New York
- Clements R, Hurst A, Knarud R, Omre H (1990) A computer program for evaluation of fluvial reservoirs. In: Buller A, Berg E, Hjemeland O, Kleppe J, Torsaeter O, Aasen J (eds) North sea oil and gas reservoirs II, pp 373–385
- Deutsch C, Cockerham P (1994) Practical considerations in the application of simulated annealing to stochastic simulation. *Math Geol* 26:67–82
- Deutsch C, Journel A (1998) GSLIB: geostatistical software library and users guide, 2nd edn. Applied geostatistics. Oxford University Press, New York
- Falivene O, Cabrera L, Muñoz J, Arbués P, Fernández O, Sáez A (2007) Statistical grid-based facies reconstruction and modelling for sedimentary bodies. Alluvial-palustrine and turbiditic examples. *Geol. Acta* 5(3):199–230
- Garboczi EJ, Day AR (1995) An algorithm for computing the effective linear elastic properties of heterogeneous materials: Three-dimensional results for composites with equal phase Poisson ratios. *J Mech Phys Solids* 43(9):1349–1503
- Geman S, Geman D (1984) Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans Pattern Anal Mach Intell* 6:721–741
- Gidas B (1989) A renormalization group approach to image processing problems. *IEEE Trans Pattern Anal Mach Intell* 11(2):164–180
- Goovaerts P (1996) Stochastic simulation of categorical variables using a classification algorithm and simulated annealing. *Math Geol* 28:909–921
- Gundeso R, Egeland O (1990) SESIMIRA; a new geological tool for 3d modelling of heterogeneous reservoirs. In Buller A, Berg E, Hjemeland O, Kleppe J, Torsaeter O, Aasen J (eds) North sea oil and gas reservoirs II, pp 363–385
- Ioannidis MA, Kwiecien MJ, Chatzis I (1997) Electrical conductivity and percolation aspects of statistically homogeneous porous media. *Transp Porous Media* 29(1):61–83
- Jeulin D (1997) Dead leaves models: From space tessellation to random functions. In Jeulin D (ed) *Advances in theory and applications of random sets*. World Scientific, Singapore, pp 137–156
- Kainourgiakis ME, Kikkinides ES, Charalambopoulou GC, Stubos AK (2003) Simulated annealing as a method for the determination of the spatial distribution of a condensable adsorbate in mesoporous materials. *Langmuir* 19(8):3333–3337
- Lehmann P, Berchtold M, Ahrenholz B, Tölke J, Kaestner A, Krafczyk M, Flüßler H, Künsch H (2008) Impact of geometrical properties on permeability and fluid phase distribution in porous media. *Adv Water Resour.* doi:[10.1016/j.advwatres.2008.01.019](https://doi.org/10.1016/j.advwatres.2008.01.019)
- Levitz P (1998) Off-lattice reconstruction of porous media: critical evaluation, geometrical confinement and molecular transport. *Adv Colloid Interface Sci* 77:71–106
- Liang Z, Ioannidis M, Chatzis I (2000) Permeability and electrical conductivity of porous media from 3d stochastic replicas of the microstructure. *Chem Eng Sci* 55(22):5247–5262
- Magnani F, Philippi P, Liang Z, Fernandes C (2000) Modelling two-phase equilibrium in three-dimensional porous microstructures. *Int J Multiph Flow* 26(1):99–123
- Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E (1953) Equation of state calculations by fast computing machines. *J Chem Phys* 21:1087–1092
- Moctezuma-Berthier A, Vizika O, Adler PM (2002) Macroscopic conductivity of vugular porous media. *Transp Porous Media* 49(3):313–359
- Moctezuma-Berthier A, Vizika O, Thovert JF, Adler PM (2004) One- and two-phase permeabilities of vugular porous media. *Transp Porous Media* 56(2):225–244
- Olayinka S, Ioannidis MA (2004) Time-dependent diffusion and surface-enhanced relaxation in stochastic replicas of porous rock. *Transp Porous Media* 54(3):273–295
- Onsager L (1944) Crystal statistics. i. a two-dimensional model with an order-disorder transition. *Phys Rev* 3(4):117–149
- Radlinski AP, Ioannidis MA, Hinde AL, Hainbuchner M, Baron M, Rauch H, Kline S (2004) Angstrom-to-millimeter characterization of sedimentary rock microstructure. *J Colloid Interface Sci* 274(2):607–612
- Roberts AP (1997) Statistical reconstruction of three-dimensional porous media from two-dimensional images. *Phys Rev E* 56(3):3203–3212

- Roberts AP, Garboczi EJ (2002) Computation of the linear elastic properties of random porous materials with a wide variety of microstructure. *Proc R Soc Lond* 458(2021):1033–1054
- Roberts AP, Knackstedt MA (1996) Structure-property correlations in model composite materials. *Phys Rev E* 54(3):2318–2328
- Roberts AP, Teubner M (1995) Transport properties of heterogeneous materials derived from Gaussian random fields: bounds and simulation. *Phys Rev E* 51(5):4141–4154
- Rozman MG, Utz M (2001) Efficient reconstruction of multiphase morphologies from correlation functions. *Phys Rev E* 63(6):066701
- Sahimi M (1993) Flow phenomena in rocks: from continuum models to fractals, percolation, cellular automata, and simulated annealing. *Rev Mod Phys* 65(4):1393–1534
- Sambridge M, Mosegaard K (2002) Monte Carlo methods in geophysical inverse problems. *Rev Geophys* 40(3). doi:[10.1029/2000RG000089](https://doi.org/10.1029/2000RG000089)
- Schonmann RH, Shlosman SB (1995) Complete analyticity for 2d Ising completed. *Commun Math Phys* 170(2):453–482
- Singh M, Mohanty KK (2000) Permeability of spatially correlated porous media. *Chem Eng Sci* 55(22):5393–5403
- Spanne P, Thovert JF, Jacquin C, Lindquist WB, Jones KW, Adler PM (1994) Synchrotron computed microtomography of porous media: topology and transports. *Phys Rev Lett* 73(14):2001–2004
- Stoyan D, Kendall WS, Mecke J (1996) *Stochastic geometry and its applications*, 2nd edn. Wiley, New York
- Szu H, Hartley R (1987) Fast simulated annealing. *Phys Lett A* 122:157–162
- Talukdar MS, Torsaeter O, Ioannidis MA (2002) Stochastic reconstruction of particulate media from two-dimensional images. *J Colloid Interface Sci* 248(2):419–428
- Torquato S (2001) *Random heterogeneous materials: microstructure and macroscopic properties*. Springer, New York
- Yeong CLY, Torquato S (1998) Reconstructing random media. *Phys Rev E* 57:495–506