

Frozen-State Hierarchical Annealing

Wesley R. Campaigne¹, Paul Fieguth¹, and Simon K. Alexander²

¹ Department of Systems Design Engineering, University of Waterloo
Waterloo, Ontario, Canada, N2L 3G1
{wrcampai, pfieguth}@uwaterloo.ca

² Department of Mathematics, University of Houston
Houston, Texas, U.S.A., 77204
simon@math.uh.edu

Abstract. There is growing demand for methods to synthesize large images of porous media. Binary porous media generally contain structures with a wide range of scales. This poses difficulties for generating accurate samples using statistical techniques such as simulated annealing. Hierarchical methods have previously been found quite effective for such problems. In this paper, a frozen-state approach to hierarchical annealing is presented that offers over an order of magnitude reduction in computational complexity versus existing hierarchical techniques. Current limitations to this approach and areas of further research are discussed.

1 Introduction

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

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

Because samples of porous media may, in some cases, be expensive to produce, handle, and measure, there is a growing literature [10,11,13] on the simulation and synthesis of large-scale 2D and 3D binary random fields. Some past approaches used a Markov-FFT approach, however the current state-of-the-art focuses on methods of statistical sampling, especially simulated annealing. Although annealing possesses attractive convergence properties *in principle*, in practice the approach is computationally very slow, particularly for images containing a mix of large-scale and small-scale structures — very common in porous media.




Fig. 1. Excerpts of two large images of physical porous media. The range of scales of pore structure (black) is clearly evident and poses challenges for pixel-based synthesis.

Our long-term research objective is the development of accelerated, hierarchical approaches to annealing-based image synthesis. We have had past success in such hierarchical synthesis, however even these approaches were hampered by having to, eventually, perform some number of annealing passes on the very finest scale. In this paper we propose a truly hierarchical method, in which large-scale structures are synthesized *and fixed in place* at coarse scales, meaning that only local, fine-resolution details remain to be refined at finer scales.

2 Motivation

Suppose we consider a two-dimensional domain of $n = m \times m$ pixels. Further suppose that the domain contains regions (solids or pore voids) having a fractional size of γ relative to the entire domain; that is, the regions have an approximate diameter $d = \gamma m$, such that γ is a fixed fraction, but that d increases as we seek to synthesize finer and finer images (larger and larger m).

Very approximately, at least $\mathcal{O}(d^2)$ passes of simulated annealing are required for convergence, thus the total computational complexity is $\mathcal{O}(\gamma^2 m^4)$. Therefore to achieve one finer scale of resolution increases the computational time by a factor of sixteen, very quickly limiting the size of image which can be considered, and therefore also limiting the range of scales which can be produced — a key consideration in porous media, many of which exhibit structure on scales many orders of magnitude apart!

Existing approaches to hierarchical simulated annealing [1,2,3] seek to ameliorate the above condition by synthesizing large-scale structure at relatively coarse scales, when the structures measure only a few pixels, allowing for much more

rapid convergence at finer scales where only small-scale details remain to be determined. Such an approach has led to two orders of magnitude improvement in computational complexity [2], however two key problems remain:

1. Because each scale is treated as a separate annealing problem, great care must be taken in selecting an annealing schedule to ensure that the structures synthesized at coarser scales are not randomly eroded.
2. Because the method continues to visit every pixel at every scale, including the finest scale, there remains a substantial computational burden of visiting very large numbers of pixels at the finest scale.

Consider the porous media such as those in Fig. 1, and suppose we wish to produce a very detailed synthesis, such as $n = 8000 \times 8000$ pixels. Clearly the overwhelming fraction of these 64-million pixels lie far within large regions, pixels which are extremely unlikely to be changed at the finest scale. That is, the overwhelming fraction of site visits at the finest scale are unnecessary. The crux of our approach is to reduce computational complexity by allowing only certain pixels to be sampled and changed at a given scale, giving rise to two benefits paralleling the above problems:

1. Because the larger-scale structures are *fixed* at coarser scales, finer scales cannot destroy or erode these structures, so our proposed method is not sensitive to a choice of annealing schedule.
2. Because our proposed method visits only a small subset of pixels at a given scale, much larger domains become computationally tractable.

3 Method

Let $x^{(s)}$ represent an image at scale s where increasing s denotes progressively finer scales, and let $x_i^{(s)}$ denote the value of $x^{(s)}$ at pixel i ; $x_i^{(s)} \in \{0, \frac{1}{2}, 1\}$ (i.e., a ternary state: black, gray, or white). Let $\{i_1, i_2, i_3, i_4\}$ be the indices of the children of $x_i^{(s-1)}$ at scale s . Then given training image $x^{(s)}$, the coarser-scale representation $x^{(s-1)}$ can be derived for modelling using

$$x_i^{(s-1)} = \begin{cases} 0 & \text{if } x_{i_j}^{(s)} = 0 \quad \forall j \\ 1 & \text{if } x_{i_j}^{(s)} = 1 \quad \forall j \\ \frac{1}{2} & \text{otherwise} \end{cases} \quad (1)$$

The repeated use of (1) allows a single sufficiently large binary image to provide data for modelling at arbitrarily coarse scales, as demonstrated in Fig. 2.

In synthesis, a corresponding set of rules is asserted as constraints on the annealing:

$$\text{If } x_i^{(s-1)} \in \{0, 1\}, \quad \text{then } x_{i_j}^{(s)} = x_i^{(s-1)} \quad (2a)$$

$$\begin{aligned} \text{If } x_i^{(s-1)} = \frac{1}{2}, \quad \text{then neither } & x_{i_j}^{(s)} = 0, \quad \forall j \\ & \text{nor } x_{i_j}^{(s)} = 1, \quad \forall j \end{aligned} \quad (2b)$$




Fig. 2. Example up-scaling and down-scaling behaviour. The heavier dividing lines indicate groups of pixels that are expressed as a single pixel at the next coarser scale, using the rule expressed in (1).

The consequence of (2) is that each state in the sample space visited by the annealer must be consistent under (1) with the result at all coarser scales. Black and white pixels represent frozen states in the image domain that are preserved at all finer scales. If a coarse-scale pixel is gray, then the corresponding set of pixels at finer scales cannot be either entirely black or entirely white. All pixels subject to (2a) can be ignored in the annealing sampler, as their values have been previously frozen.

Figure 2 demonstrates this coarsening method in practice. Progressively coarser model information is extracted via (1) based on a large binary training image. Synthesis starts from some coarse initialization, and annealing is performed at progressively finer scales (each until convergence), constrained by the previous scale’s result according to (2). When annealing the final, finest scale result, it is additionally asserted that all values must be 0 or 1.

Further constraints on x are made using some energy function $E(x)$. Choices of x leading to low energies are believed to be plausible rest states of the process. Many such energy models have been proposed, however adapting existing energy models used with binary images for use with the ternary frozen-state model is a non-trivial task. The research presented here uses an extension of the local neighbourhood model described in [1], where it was shown to be sufficient for implicitly matching a number of properties from training images.

Consider a local neighbourhood around some pixel at position i . For convenience, let us assume it is a 3×3 neighbourhood structure as in Fig. 3, with

P_1	P_2	P_3
P_4	P_0	P_5
P_6	P_7	P_8

Fig. 3. Indices for a second-order neighbourhood

position i in the central P_0 location. (Other neighbourhood structures can be used instead; this assumption is merely to simplify the explanation.)

Let c denote one of the 3^9 possible unique configurations of the neighbourhood given ternary values. The value for c , $0 \leq c < 3^9$ corresponding to the neighbourhood about pixel $x_i^{(s)}$ can be evaluated using

$$c_i^{(s)} = \sum_k (2 \cdot x_{P_k}^{(s)}) \cdot 3^k \quad (3)$$

Let $H_c^{(s)}$ then express the probability of configuration c at scale s , calculated by observation of its frequency in the training data. Figure 4 offers an example of this distribution at one scale. Let $\mathcal{H}_c(x^{(s)})$ produce histogram entry c given image $x^{(s)}$. Supposing the training data at scale s was expressed as image $y^{(s)}$, then $H_c^{(s)} = \mathcal{H}_c(y^{(s)})$. Given these terms, the energy function is then

$$E(x^{(s)}) = \sum_c (\mathcal{H}_c(x^{(s)}) - H_c^{(s)})^2 \quad (4)$$

This is a non-parametric, highly local model. To effectively capture the nature of pore structures of the sort seen in Fig. 1, this model must be paired with a hierarchical approach to image synthesis.

It was observed in early testing that at finer scales the energy function would attempt to contradict the coarse-scale results, finding it more advantageous to overfit some aspects of the training data (for instance, matching the precise probability of homogenous black or white neighbourhoods) at the expense of accurate pore shapes. For this reason, two modifications were made to the energy function.

The first modification is to weight configuration c in the energy function by its standard deviation. Let $\sigma_c^{(s)}$ express the standard deviation of the value of $H_c^{(s)}$. This is inferred by determining $H_c^{(s)}$ for many subsections of the training data at scale s with the same size as $x^{(s)}$ and evaluating the standard deviation across this set of $H_c^{(s)}$ values. Using this weighting, with some small value ϵ , a new energy function was proposed:

$$E(x^{(s)}) = \sum_c \frac{(\mathcal{H}_c(x^{(s)}) - H_c^{(s)})^2}{(\sigma_c^{(s)})^2 + \epsilon} \quad (5)$$

This modification also expresses more specifically the purpose of what the annealing is trying to accomplish: we are not trying to create an image simply




Fig. 4. A target histogram model at 64×64 . The horizontal axis indexes the 3^9 state configurations possible using the neighbourhood structure of fig. 3. The peaks at the first and last indices correspond respectively to the strong probability of homogeneously black and homogeneously white neighbourhoods.


to match one specific distribution neighbourhood *exactly*; we are trying to sample images from a distribution of plausible images. The nature of this distribution is inferred from the training data. Weighting by $\sigma_c^{(s)}$ is a means of using the variations within the training data to infer how important the precision of each value $H_c^{(s)}$ is in that greater distribution.

A second modification was to restrict the scope of the histogram operator to only those pixels whose neighbourhood configurations can be affected by the annealing process. The annealer can only affect the image as much as the constraints on the current scale permit, so the energy function should be evaluating the image *given those constraints*. Effectively, this means that $\mathcal{H}_c(x^{(s)})$ operates only on those pixels in $x^{(s)}$ that lie within the radius of the neighbourhood structure of some i_j for which $x_i^{(s-1)} = \frac{1}{2}$. This allows the energy function $E(x^{(s)})$ to evaluate the fitness of $x^{(s)}$ independently from $E(x^{(s-1)})$.

For this paper, all results came from the use of a second-order 3×3 neighbourhood, although the use of a third-order 13-pixel neighbourhood with the ternary model is also feasible. Future research may focus on adaptive methods for expanding the neighbourhood structures in a manner that actively balances the issues of memory requirements and data sparseness with increased descriptive power.

4 Results

As an initial test for the proposed algorithm, a set of artificial test images were created (Fig. 5). Each image in this set is entirely composed of a single geometric shape repeated across the image at random, non-overlapping locations with



Model input images

Synthesized results based on model

Fig. 5. Multiple images composed of equally sized shapes, with random rotations and random non-overlapping locations, were used to test the proposed algorithm’s performance in shape discrimination

random orientation. Figure 5 contains excerpts and results for the ‘circles’ and ‘triangles’ members of this set. In both cases, the proposed method, despite being non-parametric, is able to recreate the pore size and density of the respective training data. The disjointness of the pores is also preserved. The structures in the ‘circles’ image clearly resemble the circular phenomena of the training data, although the small 3×3 local neighbourhood model was unable to recreate the smooth contours at the finest scale. Similarly, for the ‘triangles’ image, the pore structure exhibits peaked points, often three per pore, but the highly local nature of the energy measurement was unable to discriminate the long, straight edges of the original image.

In both results, there are also some very small structures present — so much smaller than anything in the input images that their presence suggests an error.





Fig. 6. The number and proportion of pixels to be visited in a single pass of the annealer, as the scale increases from 8×8 to 512×512 . The number of pixels to visit at a scale equals the number of pixels which are not fixed; that is, those having a grey-valued pixel as a parent. As seen in the right plot, the pixels to visit are normally only a small fraction of the total, especially at fine scales (where the bulk of the computational effort appears) and particularly for sparsely detailed images.

They exist as a result of (2b): *some* mixture of black and white pixels must be present there to satisfy the constraint posed by a gray element at a coarser scale.






These two initial test images were designed to emulate the morphology of two physical porous media. Figure 7 is an image of a surface of a medium created using densely packed spherical glass beads. The appearance of these beads, when imaged in cross-section, is approximated in the ‘circles’ image of Fig. 5. In doing so, the ‘circles’ image tested the proposed method’s performance on this morphological aspect independently from other features present in the physical image. Similarly, the ‘triangles’ image resembles the angular, crystalline pore structures of carbonate rock in Fig. 8.

The synthesis results of Figs. 7 and 8 exhibit the same properties observed previously: the size, density, and general morphology of pore structure is maintained although the pore edges do not match the precise nature seen in the respective training images. The smaller, orphaned structures previously attributed to ternary constraints are also present, however in this instance it is not necessarily a fault: both sets of training data contain similar tiny structures. The non-ternary results, seen in Figs. 7(f) and 8(f), also exhibit such structures.

In the intermediate scales in Figs. 7 and 8, gray pixels are present on the interface between black and white regions. Gray pixels identify where structures *will* exist at finer scales; they are those pixels whose values have not yet been frozen. The progressive thinning of the gray mediation demonstrates the significance in the reduction in computational complexity, since the number of sites sampled at scale s is directly limited by the number of gray pixels present in the converged result at scale $s - 1$, plotted in Figure 6.



(a) Excerpt of training data






(b) Initial ternary constraint, 8×8 (c) Ternary result at 32×32 (d) Ternary result at 128×128 (e) Ternary result at 512×512
(full resolution)

(f) Non-ternary result

Fig. 7. Training data and synthesis results for an image of sintered glass beads surface



(a) Excerpt of training data

(b) Initial ternary constraint, 4×4 (c) Ternary result at 32×32 (d) Ternary result at 128×128 (e) Ternary result at 512×512
(full resolution)

(f) Non-ternary result

Fig. 8. Training data and synthesis results for an image of carbonate rock surface

4.1 Advantages

The most significant advantage to the proposed model is the reduction in the number of proposed changes faced by the annealer in a single pass across the image at all but the coarsest scales. The reduction experienced for the results presented is shown in Fig. 6. For the sintered glass beads results of Fig. 7, only 8.9% of the pixels in the overall hierarchy needed to be considered for changing by the annealer — an order of magnitude reduction in the computational requirements. Using the sparser carbonate rock data of Fig. 8, only a meagre 0.68% of the total pixels were examined.

The hard assertion that black and white never change at finer scales prevents large structures created at coarse scales from eroding when annealed at finer scales at high temperatures. By enabling the possibility of annealing with arbitrarily high initial temperatures at *each* scale, the annealing process is able to explore the sample space more thoroughly.

4.2 Limitations

The most immediate limitation of the frozen-state method is that other energy functions which assume a binary field, such as correlation and chordlength distributions, cannot be directly used and it is not immediately obvious how to extend such functions to work in a ternary domain.

Another consequence of using a frozen-state model is an increase in the memory needed to store the energy model. With the local neighbourhood model used here, the memory requirement for using a neighbourhood of n pixels increases as $\mathcal{O}(3^n)$, whereas it was only $\mathcal{O}(2^n)$ in the binary case. Increases to the neighbourhood size also pose the problem of increased sparseness: given limited training data, particularly at coarse scales, what is the significance of a neighbourhood configuration that shows up only once in the training set? Expanding the neighbourhood size magnifies this issue.

The ternary method is also subject to unreliable behaviour at very coarse scales. The amount of training data present at each scale is reduced by a factor of four with each level of coarseness; concerns that the data are not able to express the spatially stationary distribution of pore structures at that scale progressively increase. Consider the example in Fig. 2: Would it be accurate to assert that *each* sample from its distribution must have a single white pixel in the upper-left corner at the coarsest scale, or is that scenario merely the result of the positioning of the pore structure at the finest scale? When dealing with very limited coarse-scale data, the latter case would be a sensible assumption. Since the energy model complexity remains the same while the amount of training data decreases exponentially, the annealer is increasingly likely to overfit the energy model. This issue is amplified by the frozen-state nature of the coarse results: the hierarchical constraints asserted by the ternary model force an initial coarse overfitting to be maintained at all finer scales, reducing the diversity of synthesis results.

Although there are some limitations to the frozen-state approach, they are minor. The resistance to erosion of large structures at fine scales and, more importantly, the order-of-magnitude or greater reduction in computational complexity are of great practical benefit.

References

1. S.K. Alexander, P. Fieguth, and E.R. Vrscay, *Hierarchical annealing for random image synthesis*, EMMCVPR 2003, LNCS, no. 2683, Springer, 2003.
2. S.K. Alexander, P. Fieguth, and E.R. Vrscay, *Parameterized hierarchical annealing for scientific models*, ICIAR 2004, LNCS, no. 3211, Springer, 2004.
3. S.K. Alexander, P. Fieguth, and E.R. Vrscay, *Hierarchical Annealing for Scientific Models*, IEEE ICASSP **3** (2004), 33–36.
4. C. Bouman and M. Shapiro, *A multiscale random field model for Bayesian image segmentation*, IEEE Image Processing **3** (1994), no. 2, 162-177.
5. S. Geman and D. Geman, *Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images*, IEEE PAMI **6** (1984), 721-741.
6. B. Gidas, *A renormalization group approach to image processing problems*, IEEE PAMI **11** (1989), no. 2, 164-180.
7. T. Hofmann, J. Puzicha, and J.M. Buhmann, *Unsupervised texture segmentation in a deterministic annealing framework*, IEEE PAMI **20** (1998), no. 8, 803-818.
8. M.E. Kainourgiakis, E.S. Kikkinides, G.Ch. Charalambopoulou, and A.K. Stubos, *Simulated annealing as a method for the determination of the spatial distribution of a condensable adsorbate in mesoporous materials*, Langmuir **19** (2003), no. 8, 3333-3337.
9. Z. Kato, M. Berthod, and J. Zerubia, *A hierarchical Markov random field model and multitemperature annealing for parallel image classification*, Graphical Models and Image Processing **58** (1996), no. 1, 18-37.
10. Z. Liang, M.A. Ioannidis, and I. Chatzis, *Reconstruction of 3d porous media using simulated annealing*, Computational Methods in Water Resources XIII (Balkema, Rotterdam) (Bentley et al., ed.), 2000.
11. M.G. Rozman and M. Utz, *Efficient reconstruction of multiphase morphologies from correlation functions*, Physical Review E **63** (2001), no. 066701.
12. H. Szu and R. Hartley, *Fast simulated annealing*, Physics Letters A **122** (1987), 157-162.
13. M.S. Talukdar, O. Torsaeter, and M.A. Ioannidis, *Stochastic reconstruction of particulate media from two-dimensional images*, Journal of Colloid and Interface Science **248** (2002), no. 2, 419-428.