Parallel Hidden Hierarchical Fields for Multi-scale Reconstruction

Ying Liu and Paul Fieguth

Department of Systems Design Engineering, University of Waterloo
Waterloo, Ontario, Canada, N2L 3G1
{y30liu,pfieguth}@uwaterloo.ca

Abstract. In any problem involving images having scale-dependent structures, a key issue is the modeling of these multi-scale characteristics. Because multi-scale phenomena frequently possess nonstationary, piece-wise multi-model behaviour, the classic hidden Markov method can not perform well in modeling such complex images. In this paper we provide a new modeling approach to extend previous hierarchical methods, with multiple hidden fields, to perform reconstruction in more complex, nonstationary contexts.

1 Introduction

There are many problems in texture analysis, remote sensing and scientific imaging where the observed image possesses highly scale-dependent structure. Although such structures can, in principle, be represented with sufficiently complex models, the development of such models is a difficult task, and leads to computationally intractable algorithms if executed on a single scale. In this paper, we are interested in efficient hierarchical model structures to reconstruct complex scientific imagery.

Certainly multi-scale image modeling and analysis is common in image processing, given the widespread application of wavelet [1,2,3], Gaussian and Laplacian pyramid methods [4,5], quad-tree based models in the continuous Gaussian case [6, 7], and texture synthesis [8]. In addition, multi-fractal analysis is used to characterize the self-similarity property of objects. This technique has been used to study the statistics of natural images [9], been applied to synthesize textures [9], and acted as a prior to regularize reconstruction problems [10]. However, these methods are all applied to continuous-state problems, whereas our interest is in the hierarchical modeling of discrete-state fields, such as a label field underlying an image.

In the discrete-state case, Markov / Gibbs Random Fields have been widely used in image restoration, segmentation, reconstruction [11, 12]. However, local Markov / Gibbs Random Fields are limited to describing phenomena at a single scale, and are not naturally suited for multi-scale phenomena.

Instead, hierarchical MRF modeling [13] provides a more natural and efficient way to deal with multi-scale structures. Kato et al [14] proposed a hierarchical

D. Cremers et al. (Eds.): EMMCVPR 2009, LNCS 5681, pp. 70–83, 2009.

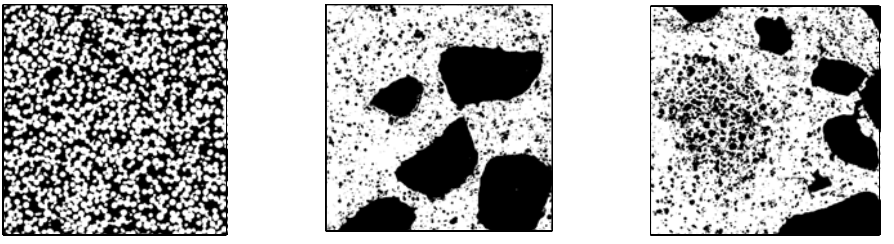

(a) Single-scale example (b) Two-scale example (c) Complex, multi-scale, multi-model

Fig. 1. Excerpts from microscopic images of physical porous media. A single-scale structure (a) can be well described by a single hierarchical MRF [17], two-scale nonstationary behaviours in (b) may be described by adding a hidden hierarchical MRF [18], however complex structures (c) with multi-model behaviours pose a modeling challenge.

MRF model with a 3D neighborhood system for modeling label fields, but with considerable computational cost. A MRF model based on a quad-tree structure was discussed by Laferté et al [15], but does not model the interactions within scales. Later, Mignotte et al [16] proposed a model with an inter-scale Markov chain and an intra-scale MRF. However, all of these used simple models at each scale, too limited to capture complex structures, such as those in Fig. 1.

In order to precisely synthesize images of porous media, Alexander and Fieguth [19] proposed a hierarchical model with local MRFs at each scale, but ignoring interrelations between scales. Later, Campaigne [17] proposed a frozen-state hierarchical annealing method, with attractive computational complexity and scale-dependent modeling.

The goal of our research is the extension of [17] to allow hidden fields. Generally, a single hierarchy with a scale-dependent model can capture a stationary structure (Fig. 1(a)), whereas many random fields have some sort of nonstationary piece-wise multi-model behaviour which requires additional hidden fields [20, 21, 22, 23]. Although multiple hidden fields are routinely used in Markov modeling, asserting a hierarchical context creates additional subtleties. Recently, Scarpa et al [24] proposed a hierarchical texture model which represents texture at the region level with a superimposed finite-state hierarchical model. Their approach has some similarities with ours, but focuses on unsupervised model inference, whereas our approach requires more accurate, supervised models, with an emphasis on computational tractability for large problems.

In this paper we build on previous work [18] to perform reconstruction for complex, nonstationary problems. We have chosen to apply our methods to reconstruct scientific images from porous media, such as the one shown in Fig. 1(c), since the images include multiple behaviours, with fractal-like scale dependent structures. The problem reduces to an energy minimization across fields and across scales, with promising results.

2 Background

2.1 Classical Hidden Markov Framework

Based on the hidden Markov field (HMF) [25,26] framework, image reconstruction can be achieved by estimating a hidden random field X from an observed field Y, where $Y = \{Y_s : s \in S_L\}$ is defined on a Low Resolution (LR) grid space S_L, and $X = \{X_s : s \in S_H\}$ is defined on a High Resolution (HR) grid space S_H with a size of $N \times N$. The relationship between X and Y is expressed by a forward model $Y = g(X) + \nu$, where ν denotes the measurement noise. If ν is i.i.d, the classical HMF is written as

$$p(x|y) \propto \prod_{s \in S_L} p(y_s|x)p(x) \tag{1}$$

where X is assumed to be MRF. However, a single local MRF cannot perform well in modeling a multi-scale nonstationary X; for example Fig. 3(c) shows the failure of the classical HMF method to reconstruct a piece-wise, two-scale image (Fig. 3(a)).

2.2 Single Hierarchical Field

In modeling multi-scale phenomena, the dilemma of a single MRF is that local models cannot strongly assert the presence of nonlocal structures, whereas learning a huge nonlocal model is prohibitive in practice. Very differently, hierarchical modeling defines X via a sequence of fields $\{X^k,\, k \in K = (0, 1, \cdots, M)\}$, where $k = 0$ defines the finest scale. At each scale k, X^k is defined on site space S^k and results from the downsampling of X^0.

The advantage of hierarchical modeling is that nonlocal large-scale features become local at a sufficiently coarse scale, therefore at each scale a single MRF (X^k) can be used to capture the features local to that scale, inherently allowing scale-dependent structure. In defining a hierarchical model, two issues need emphasizing: the inter-scale context, and the computational complexity.

To model the spatial context, Mignotte et al [16] proposed a Markov chain in scale $p(x^k|x^{K\setminus k}) = p(x^k|x^{k+1})$, whereas the intra-scale relation is a MRF $p(x_s^k|x_{S\setminus s}^k) = p(x_s^k|x_{\wp(s)}^{k+1}, x_{\mathcal{N}_s^k}^k)$, where $\wp(s)$ denotes the parent site of s at the parent scale, and $\mathcal{N}_s^k$ defines a local neighborhood.

To achieve computational efficiency, a Frozen State Hierarchical Field (FSHF) was presented in [17] to synthesize binary images. In their work, a given HR field ($x = x^0$) can be represented by a hierarchical field $\{x^k\}$ (Fig. 2) where $x^k = \Downarrow^k(x^0)$, and $\Downarrow^k(\cdot)$ denotes a downsampling operator. The key idea of the FSHF is that, at each scale, only those sites which are undetermined need to be sampled, with those sites determined by the parent scale fixed (or frozen):

$$p(x_s^k|x_{S\setminus s}^k) = \begin{cases} \delta_{x_s^k, x_{\wp(s)}^{k+1}}, & \text{if } x_{\wp(s)}^{k+1} \in \{0,1\} \\ p(x_s^k|x_{\mathcal{N}_s^k}^k), & \text{if } x_{\wp(s)}^{k+1} = \frac{1}{2} \end{cases} \tag{2}$$

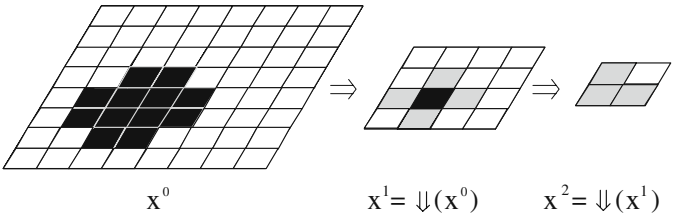

Fig. 2. An example of ternary hierarchical subsampling [17]: a given field x^0 is coarsified by repeated 2×2 subsampling $\Downarrow(\cdot)$. All-white and all-black regions are preserved, with mixtures labeled as uncertain (grey).

where $0, 1$ (black, white) are determined states, and $\frac{1}{2}$ (grey) is undetermined. Since the "grey" interface between black and white regions represents only a small fraction of most images, this approach offers a huge reduction in computational complexity relative to standard, full-sampling hierarchical techniques. The site sampling strategy corresponding to (2) is

$$\hat{x}_s^k = \begin{cases} \hat{x}_{\wp(s)}^{k+1} & \text{if } \hat{x}_{\wp(s)}^{k+1} \in \{0,1\} \\ \\ \text{a sample from } \; p(x_s^k | x_{\mathcal{N}_s^k}^k) & \text{if } \hat{x}_{\wp(s)}^{k+1} = \frac{1}{2} \end{cases} \tag{3}$$

where $\hat{x}^k$ is the sampled (estimated) random field at scale k.

2.3 Hidden Hierarchical Fields

A single FSHF works well in modeling stationary scale-dependent structures, however such a model cannot handle X having nonstationary, piece-wise behaviour, because conditioned on X^{k+1}, X^k still has nonstationary features which cannot be captured by a local model. For example, Fig. 3(d) shows that a single hierarchy can not capture the piece-wise two-model behaviour in Fig. 3(a).

To model more general multi-scale cases, we proposed a Hidden Hierarchical Markov Field (HHMF) in [18]. The HHMF has an hidden binary HR field U to capture the model behavior in X. If the nonstationarity in X can be entirely attributed to a single binary behaviour, then conditioned on U becomes X conditional stationary. Assuming X, U to be Markov in scale, the joint relationship $p(X, U) = p(X|U)p(U)$ can be written as

$$p(x, u) \propto \left[\prod p(x^k | x^{k+1}, u^k)\right] \left[\prod p(u^k | u^{k+1})\right] \tag{4}$$

We select some coarsest scale as the scale at which decidable state structure appears, such that (4) becomes

$$p(x, u) \propto \left[\prod_{k=0}^{M_x - 1} p(x^k | x^{k+1}, u^k)\right] p(x^{M_x} | u^{M_x}) \left[\prod_{k=0}^{M_u - 1} p(u^k | u^{k+1})\right] p(u^{M_u}) \tag{5}$$

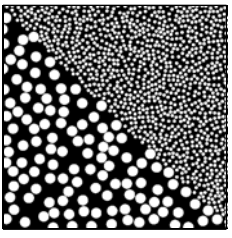

(a) Ground truth $\overset{*}{x}$, 512×512

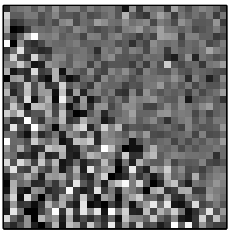

(b) LR, noisy y, 32×32

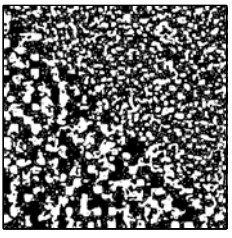

(c) HR estimate $\hat{x}$ based on single MRF,512×512

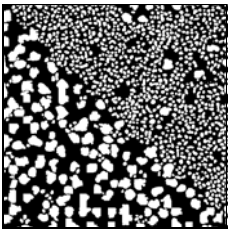

(d) HR estimate $\hat{x}$ based on FSMF, 512×512

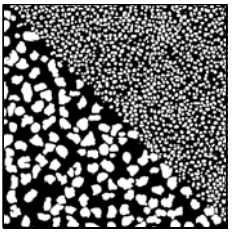

(e) HR estimate $\hat{x}$ based on HHMF, 512×512

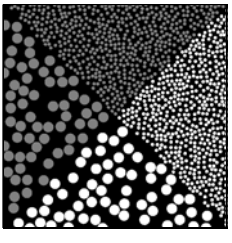

(f) HR image with two spatial variables

Fig. 3. In this example, we reconstruct a two-scale image (a) from a low resolution measurement (b) with different Hidden Markov Field frameworks. The estimated results are shown in (c)-(e). The clear scale separation of the result from the hidden hierarchy [18] (e) should be compared to the results from single flat MRF based model [27] (c) and single hierarchy based model [17] (d). (f) shows an image with two spatial variables, which creates a modeling challenge for existing models.

where $k = M_x$ and $k = M_u$ denote the coarsest scale of X and U respectively, and where $p(x^k|x^{k+1}, u^k)$ and $p(u^k|u^{k+1})$ are modeled as frozen-state (2), based on the joint, local ternary histogram of [19]. Since U is defined to describe large scale features or model behaviour in X, the decidable state in X is expected to vanish at a finer scale than in U ($M_x < M_u$).

The introduction of a hierarchical hidden field provides a more general, powerful modeling ability, as illustrated in the comparison between Fig. 3(b) and (e), however more complex problems (Fig. 3(f) or Fig. 1(c)) require a more general approach, the subject of this paper.

3 Parallel Hidden Hierarchical Fields

In general, the behaviour of a random field X will be determined by more than one spatial variable ($N_v > 1$), such that X remains nonstationary when

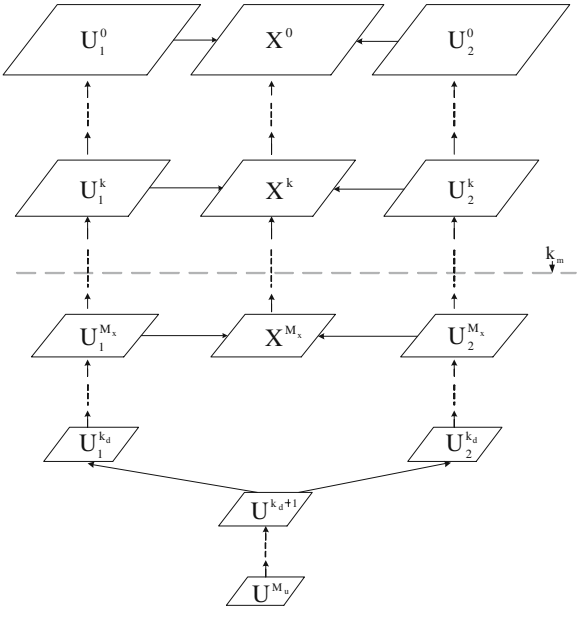

Fig. 4. The proposed Parallel Hidden Hierarchical Markov Field modeling structure, such that the hidden label field U is a joint field only at coarse scales. As the features of different model behaviour become separable at some scale k_d, U is decoupled as multiple parallel hierarchical fields $\{U_1, \cdots, U_{N_v}\}$. At scales coarser than M_x the entire random field X of interest is uncertain, and so only U is represented to scale M_u.

conditioned on a single binary field U. The obvious solution to this problem is to define U as a multi-label field; for example the behaviour of Fig. 3(f) is determined by two variables of scale and shade, with a corresponding quad-label model.

Generalizing the frozen-state annealing algorithm [17] to the non-binary case is a non-trivial task, yet the significant computational benefits of the frozen-state approach motivate us to continue using it. Although multi-label models (e.g., Potts) do exist, the complexity of modeling all pairwise, triplet-wise etc. label interactions at coarser scales makes the problem rather complex.

Our proposed approach is to use multiple, parallel hierarchical hidden fields (PHHF), as illustrated in Fig 4. At finer scales, having many state elements, the hidden fields are decoupled, binary, and simply modeled. The complex, joint hidden structure appears only at very coarse scales, where the small number of state elements allows such a structure to be computationally tractable.

3.1 Parallel Hidden Hierarchical Field

The key idea of PHHF is as follows: instead of using a single hidden field U to model complex multi-model behaviors directly, we introduce multiple binary hidden label fields $U = \{U_i, i \in N_v\}$, such that each field U_i is used to capture

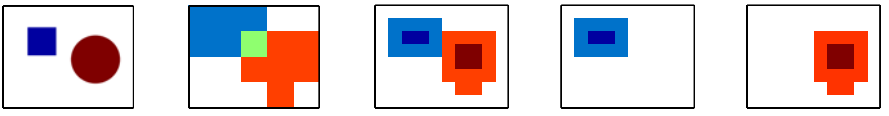

(a)Label Field u (b) u^7 (c) u^6 (d) u_1^6 (e) u_2^6

Fig. 5. For a hidden label field U with nonlocal features (a), at some coarse scale the features of different behaviours interact (b), however at some finer scale the U become separable (c), and can be decoupled to multiple simple fields (d,e)

only a single binary structure or model behavior. With this assumption, the joint field, conditioned on its parent, decouples into independent fields

$$p(u_1^k, \cdots, u_{N_v}^k | u^{k+1}) = \prod_{i=1}^{N_v} p(u_i^k | u^{k+1}) \tag{6}$$

for all scales $k \leq k_d$. The entire PHHF model can then be extended from (5) and written as

$$p(x,u) \propto \prod_{k=0}^{M_x-1} p(x^k|x^{k+1},u^k) \cdot p(x^{M_x}|u^k) \cdot \prod_{k=0}^{k_d-1} \prod_{i=1}^{N_v} p(u_i^k|u_i^{k+1}) \cdot \prod_{i=1}^{N_v} p(u_i^{k_d}|u^{k_d+1}) \prod_{k=k_d+1}^{M_u-1} p(u^k|u^{k+1}) \cdot p(u^{M_u}) \tag{7}$$

An example of the PHHF with two hidden variables ($N_v = 2$) is shown in Fig. 4. The approach simplifies modeling in three significant ways.

First, the PHHF consists entirely of simple models, both local and stationary. Specifically, although X^k and U^k may have complex, non-local behaviour, the conditional residuals $U^k|U^{k+1}$, $X^k|X^{k+1}, U^k$ are local, by virtue of the fact that all non-local matters have been absorbed into the conditioned (coarser) scale.

Second, the hidden states are primarily decoupled and binary. At coarse scales, where only few pixels exist, it is computationally tolerable to assert a joint model for $U^k, k > k_d$, where the joint model is need to allow the hidden models to interact (Fig. 5(b)). In most problems, empirically, the hidden models become separable at some scale $k \leq k_d$ (Fig. 5(c)), leading to parallel independent fields (Fig. 5(d)(e)).

Third, because $\{X^k\}$ and $\{U_i^k\}$ are modeled using simple, binary models, $\{X^k\}$ and $\{U_i^k\}$ are easily defined as hierarchical frozen states, leading to the computational cost of the PHHF being linear in the number of hidden fields N_v, except at the small, coarse scales.

3.2 Reconstruction

To reconstruct a scale-dependent, near fractal, piece-wise nonstationary image such as the porous medium in Fig. 1(c) is a major modeling challenge. The image

in Fig. 1(c) displays three types of behaviour: large-scale pores, regions of high density, and background areas of low density. We therefore propose the ternary hidden field U to be decoupled as two parallel hierarchies, where $\{U_1^k\}$ identifies large pores, and $\{U_2^k\}$ identifies regions of high density.

For image enhancement, the hidden fields are invisible to the measurements, therefore $p(y|x^k, u^k) = p(y|x^k)$, and so the reconstruction model illustrated in Fig. 4 can be written as

$$p(x, u|y) \propto \prod_{s \in S_L} p(y_s|x)p(x, u) \tag{8}$$

where $p(x, u)$ is the prior PHHF, defined as in (7), and where the measurements $p(y_s|x)$ are taken at some scale k_m.

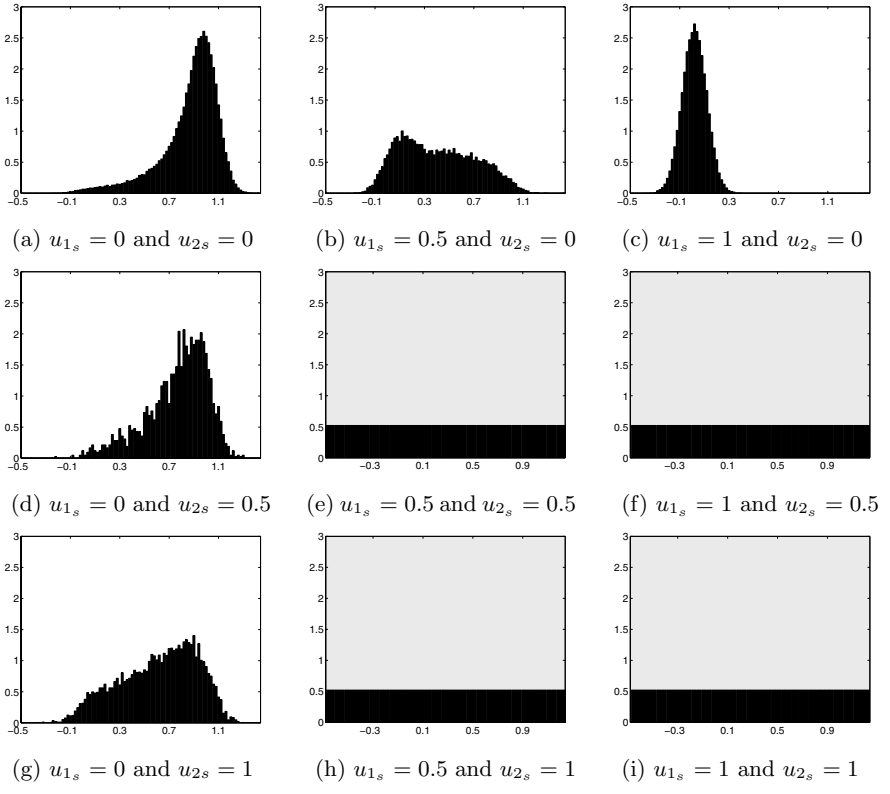

(a) $u_{1_s} = 0$ and $u_{2s} = 0$ (b) $u_{1_s} = 0.5$ and $u_{2_s} = 0$ (c) $u_{1_s} = 1$ and $u_{2_s} = 0$

(d) $u_{1_s} = 0$ and $u_{2s} = 0.5$ (e) $u_{1_s} = 0.5$ and $u_{2_s} = 0.5$ (f) $u_{1_s} = 1$ and $u_{2_s} = 0.5$

(g) $u_{1_s} = 0$ and $u_{2s} = 1$ (h) $u_{1_s} = 0.5$ and $u_{2_s} = 1$ (i) $u_{1_s} = 1$ and $u_{2_s} = 1$

Fig. 6. The conditional target histograms of $g(X)$ for decoupling a joint field U^{k_d} into two simple fields $U_1^{k_d}$ and $U_2^{k_d}$; that is, each panel shows the distribution of $g(X)$ for one of nine possibilities on U_1, U_2. Since the hidden fields are asserted to be decoupled, those cases where both fields are asserted (shaded distributions) are never observed, and so are assigned a uniform distribution with low marginal probability. To the extent that the joint state configuration of $(U_{1_s}^{k_d}, U_{2_s}^{k_d})$ relates to distinguishable model behaviour in $g(X)$, we expect the hidden fields to be estimable.

Given measurements Y contaminated by i.i.d. noise, the posterior distribution of (X, U, Y) can be represented as a Gibbs distribution

$$p(x, u|y) \propto exp(-\frac{1}{T}E(x, u|y)) \tag{9}$$

where T is the temperature, such that E is the energy function implying the probability density p. Finding a good estimate $\hat{x}$ therefore corresponds to maximizing $p(x, u|y)$, correspondingly minimizing E, the sum of hidden joint $E_u^k(u^k|u^{k+1})$, decoupled $E_{u_i}^k(u_i^k|u_i^{k+1})$, visible $E_{x|u}^k(x^k|x^{k+1}, u^k)$, and measurement $E_m(y|x)$.

All of the prior models are learned using a nonparametric joint local distribution, the exhaustive joint distribution of a local 3×3 neighbourhood of ternary state elements. The models are learned separately on each scale, based on downsampled training data $\bar{x}^k = \Downarrow^k (\bar{x}^0|\bar{u}^0)$, $\bar{u_i}^k = \Downarrow^k (\bar{u_i}^0)$. The resulting energy function is the least-squares difference between the model and observed joint histograms [19].

The measurement energy function is inferred from the given forward model $g()$. A variety of measurements could be defined, depending on the measuring instrument, however in this paper we focus on reconstruction from low-resolution images, making $g()$ a downsampling operator.

To minimize $E(x, u|y)$, we need to anneal on each scale k in each field X, U_i, with consequent open questions: whether to minimize hidden states separately or jointly with the observable state, whether to minimize the scales in parallel or sequentially, and whether to have scale-dependent annealing schedules. Definitive answers to these questions are unknown, however empirical testing suggests that a constant annealing schedule with sequential minimization over scales and sequential minimization from hidden (U) to visible states (X) lead to a reliable and robust reconstruction.

When estimating the hidden field U, in which case X is unknown, the distribution $p(y|u)$ is less obvious, and needs to be inferred empirically, as illustrated in Fig. 6, from downsampled ground-truth of the hidden fields. Fig. 6 plots the

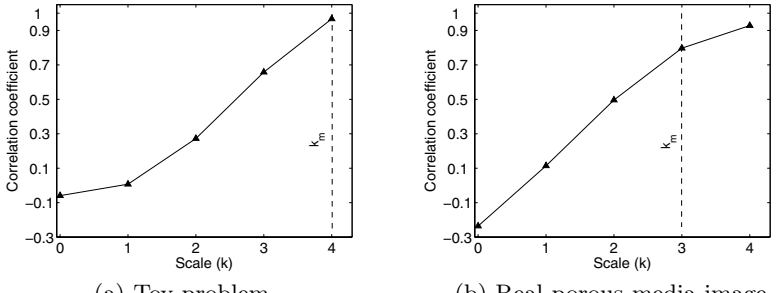

(a) Toy problem (b) Real porous media image

Fig. 7. Correlation coefficients ρ between the estimates $\hat{x}$ and ground truth $\overset{*}{x}$ as a function of structure scale. For a number of scales below the measured resolution k_m, $\rho(\overset{*}{x}, \hat{x}) > 0$ meaning that some trustable details are reconstructed.

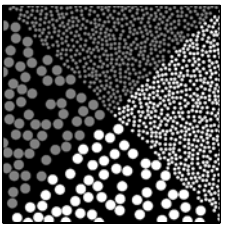

(a) Ground truth $\overset{*}{x}$, 512×512

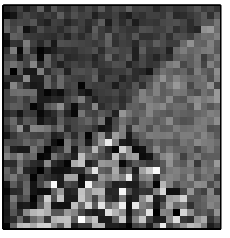

(b) LR, noisy y, 32×32

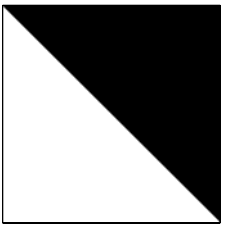

(c) Hidden scale field $\overset{*}{u}_1$, 512×512

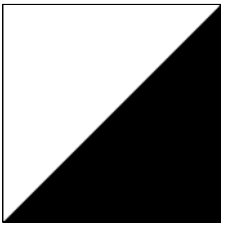

(d) Hidden shade field $\overset{*}{u}_2$, 512×512

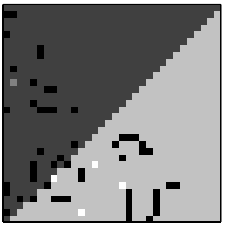

(e) $\hat{x}^4$, 32×32

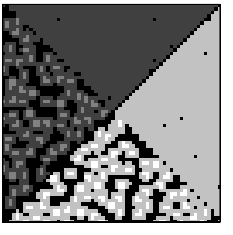

(f) $\hat{x}^3$, 64×64

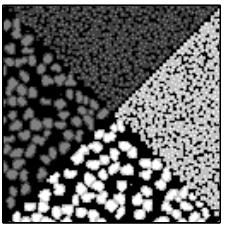

(g) $\hat{x}^2$, 128×128

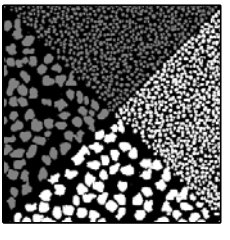

(h) $\hat{x}^1$, 256×256

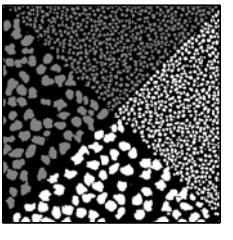

(i) HR estimate $\hat{x}$ from PHHF, 512×512

Fig. 8. A toy problem with two spatial variables. For the purpose of this example, we assume the hidden fields for scale and shade to be given, in (c) and (d) respectively. From the low resolution measurements of (b), our estimated results $\{\hat{x}^k, 0 \leq k \leq 4\}$ are shown in (e)-(i). A clear scale and shade separation are shown in the final result (i), demonstrating the strength of modeling by paralleled hidden fields.

nonparametric histogram in Y as a function of the nine possible joint relationships in U_1 and U_2. Because the hidden fields are decoupled, four of the nine joint relationships are inadmissible (shown as shaded, in the figure), and are modeled as uniform, with a low marginal probability.

The sequential estimation from U to X has the further benefit of determining U in detail, at the pixel level, when estimating X, avoiding the ambiguity of an "undecided" (state $\frac{1}{2}$) hidden state with ambiguous assertions on X.

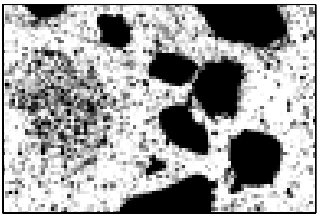

(a) LR downsample, $g(\overset{*}{x})$, 64×96

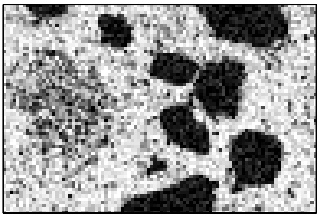

(b) Measurement, $y = g(\overset{*}{x}) + \omega$, 64×96

(c) True label field of $\overset{*}{u}_1$, 512×768

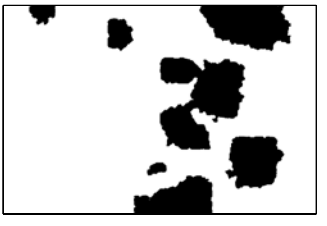

(d) Estimated label field,$\hat{u}_1$ 512×768

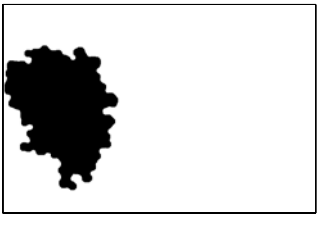

(e) True label field of $\overset{*}{u}_2$, 512×768

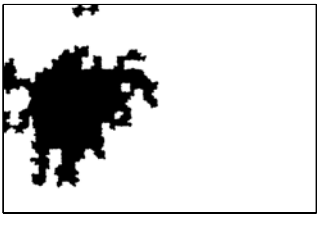

(f) Estimated label field,$\hat{u}_2$ 512×768

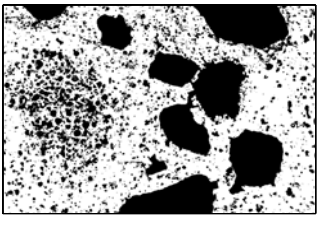

(g) Ground truth, $\overset{*}{x}$, 512×768

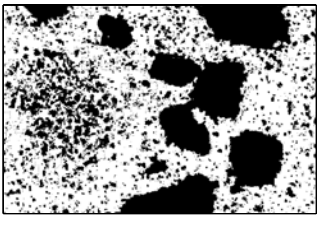

(h)HR estimate, $\hat{x}$, 512×768

Fig. 9. Ground truth and reconstruction results for a real porous media image. Although $\overset{*}{x}$ is not able to fully reconstruct some subtle structures (eg., the connectivities at the interface between the large pores), the improvement in relevant detail of (h) over (b) is stunning. Here, $k_m = 3$, $M_x = 3$, $k_d = 5$, and $M_u = 6$.

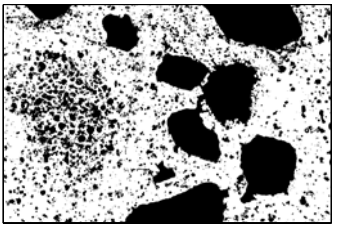

(a) Ground truth, $\overset{*}{x}$, 512×768

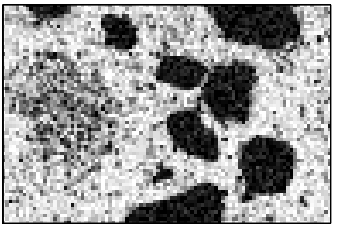

(b) Measurement, $y = g(\overset{*}{x}) + \omega$, 64×96

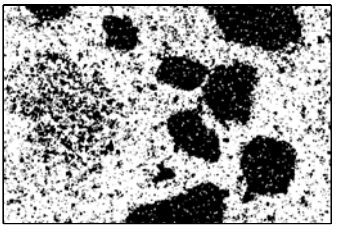

(c) HR estimate $\hat{x}$ from the classical HMF [27], 512×768

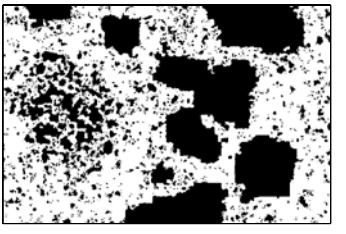

(d) HR estimate $\hat{x}$ with single FSMF [17], 512×768

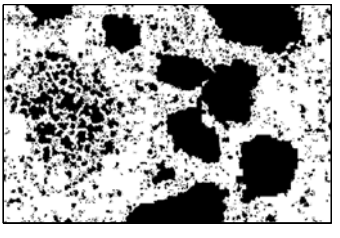

(e) HR estimate $\hat{x}$ with, HHMF [18], 512×768

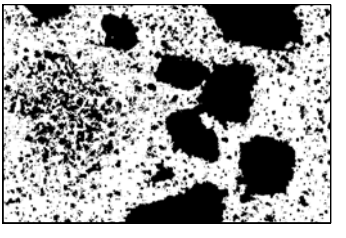

(f) HR estimate, $\hat{x}$ with the proposed method, 512×768

Fig. 10. Reconstruct results for a real porous media image from different frameworks. Compared to the results from other methods (c)-(e), the reconstruction result $\hat{x}$ from our proposed method (f) shows an impressive improvement in modeling multiple spatial nonstationarities. In particular, observe closely the small-scale structures in the high dense regions and the boundaries of large pores.

The computational cost of our proposed approach can be approximated as

$$C_{PHHA} \simeq \sum_{k=0}^{M_u} (N_\nu + 1) \cdot \alpha_k \cdot (4^{-k} N^2) \tag{10}$$

where α_k denotes the fraction of unfrozen pixels at each scale k. Experimentally, C_{PHHA} remains dominated by the cost of the finest scale $O((N_\nu + 1) \cdot \alpha_0 \cdot N^2)$.

4 Results and Conclusions

Our results will be based on a synthetic four-label image (Fig. 3(f)) and a real porous media image with multi-scale, complex structure (Fig. 1(c)).

In the synthetic example we have two hidden fields U_1, U_2, which respectively describe the states of circle-size and shade. For the purpose of estimating X, we will assume U_1, U_2 to be known (Fig. 8(c,d)). Given a LR noisy image (Fig. 8(b)), our estimated results $\{\hat{x}^k, 0 \leq k \leq 4\}$ are shown in (Fig. 8(e)-(i)). Clearly, in the reconstruction process the structures of the two-scale circles are gradually decided from coarse to fine with different label values. The impressive reconstruction result (Fig. 8(i)) illustrates the positive effect of the two parallel fields U_1 and U_2 to label the nonstationary behaviours.

In the porous media example, a much more difficult problem, we estimate first U_1, U_2 and then X, as proposed. The final HR estimate ($\hat{x}$) is shown in Fig. 9. The performance of the proposed PHHF is clear from the comparison to other methods (Fig. 10(c-e)).

Finally, to demonstrate the the reconstruction is actually correctly estimating fine-scale details, Fig. 7 plots the correlation coefficient between ground truth and the reconstruction $\hat{x}^0$ as a function of structure scale (defined as the average number of decimations which leaves a pixel value unchanged). Clearly strongly positive correlations exist well below the measured scale, meaning that the enhancement is adding relevant detail, and not just synthesizing random structure from the prior model. Clearly very tiny structures fail to exert much influence on the measurements, therefore the correlation ρ does decrease at finer scales.

Our research interest is the hierarchical representation of multi-label hidden fields. Although here the proposed approach is applied only to porous media images, it can be extended to more general problems in modeling, analysis, and processing. Our intention is to move this research towards hierarchical multi-label synthesis, and the reconstruction of three-dimensional data from low-resolution observations.

References

1. Mallat, S.G.: A theory for multiresolution signal decomposition. IEEE Trans. Pattern Anal. Mach. Intell. 11, 674–693 (1989)
2. Crouse, M.S., Nowak, R.D., Baraniuk, R.G.: Wavelet-based statistical signal processing using hidden markov models. IEEE Trans. Signal Proc. 46, 886–902 (1998)
3. Choi, H., Baraniuk, R.: Multiscale image segmentation using wavelet-domain hidden markov models. IEEE Trans. Image Processing 10, 1309–1321 (2001)
4. Burt, B., Adelson, E.: The laplacian pyramid as a compact image code. IEEE Trans. Commun. 31, 532–540 (1983)
5. Olkkonen, H., Pesola, P.: Gaussian pyramid wavelet transform for multiresolution analysis of images. Graphical Models and Image Proceesing 58(32), 394–398 (1996)
6. Basseville, M., Benveniste, A., Chou, K.C., Golden, S.A., Nikoukhah, R., Willsky, A.S.: Modeling and estimation of multiresolution stochastic processes. IEEE Trans. Inform. Theory 38, 766–784 (1992)

7. Fieguth, P.: Hierarchical posterior sampling for images and random fields. IEEE Trans. Image Processing (accepted, 2008)
8. Han, C., Risser, E., Ramamoorthi, R., Grinspun, E.: Multiscale texture synthesis. ACM Trans. Graph 27(3), 51:1–51:8 (2008)
9. Chainais, P.: Infinitely divisible cascades to model the statistics of natural images. IEEE Trans. Pattern Anal. Mach. Intell. 29, 2015–2119 (2007)
10. Liao, I.Y., Petrou, M., Zhao, R.: A fractal-based relaxation algorithm for shape from terrain image. Computer Vision and Image Understanding 109, 227–243 (2008)
11. Li, S.: Markov Random Field Modeling in Image Analysis. Springer, Heidelberg (2001)
12. Chellappa, R., Jain, A.: Markov Random Fields: Theory and Applications. Academic Press, London (1993)
13. Graffigne, C., Heitz, F., Perez, P.: Hierchical markov random field models applied to image analysis: a review. In: SPIE, vol. 2568, pp. 2–17 (1994)
14. Kato, Z., Berthod, M., Zerubia, J.: A hierarchical markov random field model and multitemperature annealing for parallel image classificaion. Graphical Models and Image Proceesing 58(1), 18–37 (1996)
15. Laferté, J.M., Pérez, P., Heiz, F.: Discrete markov image modeling and inference on the quadtree. IEEE Trans. Image Processing 9(3), 390–404 (2001)
16. Mignotte, M., Collet, C., Pérez, P., Bouthemy, P.: Sonar image segmentation using an unsupervised hierarchical mrf model. IEEE Trans. Image Processing 9(7), 1216–1231 (2000)
17. Campaigne, W., Fieguth, P., Alexander, S.: Forzen-state hierarchical annealing. In: Campilho, A., Kamel, M.S. (eds.) ICIAR 2006. LNCS, vol. 4141, pp. 41–52. Springer, Heidelberg (2006)
18. Liu, Y., Fieguth, P.: Image resolution enhancement with hierarchical hidden fields. In: ICIAR 2009 (2009)
19. Alexander, S.K., Fieguth, P., Vrscay, E.R.: Hierarchical annealing for random image synthesis. In: Rangarajan, A., Figueiredo, M.A.T., Zerubia, J. (eds.) EMMCVPR 2003. LNCS, vol. 2683. Springer, Heidelberg (2003)
20. Benboudjema, D., Pieczynski, W.: Unsupervised statistical segmentation of nonstationary images using triplet markov fields. IEEE Trans. Pattern Anal. Mach. Intell. 29(8), 1367–1378 (2007)
21. Marroquin, J.L., Santana, E.A., Botello, S.: Hidden markov measure field models for image segmentation. IEEE Trans. Pattern Anal. Mach. Intell. 25(11), 1380–1387 (2003)
22. Rivera, M., Gee, J.C.: Twoclevel mrf models for image restoration and segmentation. In: BMVC 2004, pp. 809–818 (2004)
23. Liu, Y., Mohebi, A., Fieguth, P.: Modeling of multiscale porous media using multiple markov random fields. 4th Biot (June 2009)
24. Scarpa, G., Haindl, M., Zerubia, J.: A hierarchical texture model for unsupervised segmentation of remotely sensed images. In: Ersbøll, B.K., Pedersen, K.S. (eds.) SCIA 2007. LNCS, vol. 4522, pp. 303–312. Springer, Heidelberg (2007)
25. Geman, S., Geman, D.: Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. IEEE Trans. Pattern Anal. Mach. Intell. 6(6), 721–741 (1984)
26. Besag, J.: On the statistical analysis of dirty pictures. J. Roy. Statist. Soc. Ser. B 48(3), 256–302 (1986)
27. Mohebi, A., Fieguth, P.: Posterior sampling of scientific images. In: Campilho, A., Kamel, M.S. (eds.) ICIAR 2006. LNCS, vol. 4141, pp. 339–350. Springer, Heidelberg (2006)