

Published as:

B. Wohlberg and G. de Jager, "A Class of Multiresolution Stochastic Models Generating Self-Affine Images," *IEEE Transactions on Signal Processing*, vol. 47, no. 6, pp. 1739-1742, June 1999.

IEEE Copyright Notice

©1999 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

This material is presented to ensure timely dissemination of scholarly and technical work. Copyright and all rights therein are retained by authors or by other copyright holders. All persons copying this information are expected to adhere to the terms and constraints invoked by each author's copyright. In most cases, these works may not be reposted without the explicit permission of the copyright holder.

A Class of Multiresolution Stochastic Models Generating Self-Affine Images

Brendt Wohlberg and Gerhard de Jager

Abstract— Fractal image compression is based on the rather poorly motivated assumption that “natural” images exhibit significant affine self-similarity. The accuracy of this assumption is evaluated by a comparison between the statistics of natural images and those of a multiresolution stochastic model designed to generate images exhibiting affine self-similarity as assumed by fractal coding techniques. These comparisons suggest that self-affinity does not represent a particularly accurate characterisation of image statistics.

Keywords— Fractals, Higher order statistics, Image coding, Stochastic image models, Wavelet transforms

I. INTRODUCTION

Fractal coding involves the representation of an image by a contractive affine transform of which the fixed point is close to the original image [1]. While the original fractal coding schemes constructed this transform as the composition of mappings between image subblocks, analysis is simplified by considering this construction in the wavelet transform domain, where the mappings between subblocks correspond, given some restrictions on the fractal coding scheme, to mappings between subtrees of detail coefficients [2]. In the simplest form of coding within this framework, all detail coefficients at some fixed resolution (i_d) become the roots of “domain” subtrees, the detail coefficients at the next resolution ($i_r = i_d + 1$) forming the roots of the “range” subtrees (see Figure 1). An image is encoded by finding a matching domain subtree for each range subtree, such that some scalar multiple of the domain is sufficiently close to the range; the fractal representation consists of all detail coefficients up to and including resolution i_d , together with the identity of the domain subtree and associated scaling factor¹ selected for each range subtree.

Decoding is best explained by reference to Figure 1, in which $i_d = 1$. Assume that D_0 is the matching domain for R_0 and R_2 , and D_1 is the matching domain for R_1 and R_3 , and label the scaling factor for the domain mapped to R_j as s_j . Each range subtree is reconstructed by applying the corresponding mapping defined during the encoding procedure - the matching domain subtree is scaled and copied onto the range subtree. Since detail coefficients at resolutions greater than i_d are not available when decoding is initiated, a single level of the subtree mappings is applied to generate the detail coefficients at resolution

B. Wohlberg was with the Digital Image Processing Laboratory, Electrical Engineering Department, University of Cape Town, Rondebosch 7701, South Africa. He is now with Los Alamos National Laboratory, Los Alamos, NM 87545, USA. G. de Jager is with the Digital Image Processing Laboratory, Electrical Engineering Department, University of Cape Town, Rondebosch 7701, South Africa.

¹These scaling factors should not be confused with the scaling coefficients associated with the scaling function in wavelet analysis.




Fig. 1. Domain (D_0, D_1) and range (R_0, R_1, R_2, R_3) subtrees in the wavelet decomposition of a one-dimensional signal. Note that the highest resolution detail coefficients are omitted from the domain subtrees to ensure that they have the same size as the range subtrees.

$i_d + 1$ ($d_{2,0} = s_0 d_{1,0}$, $d_{2,1} = s_1 d_{1,1}$, $d_{2,2} = s_2 d_{1,0}$, and $d_{2,3} = s_3 d_{1,1}$ in Figure 1), the next level of mappings generates resolution $i_d + 2$ ($d_{3,0} = s_0 d_{2,0}$, $d_{3,1} = s_0 d_{2,1}$, etc. in Figure 1), and so on until all necessary detail coefficients have been generated.

If isometry operations in the spatial domain are avoided, there is no mixing between coefficients in the horizontal, vertical and diagonal subbands of a non-standard two-dimensional wavelet basis [2], and one may consider decoding separately in each of these three families of subbands; encoding based on combined subtrees from all three directions merely implies that domain positions and associated scaling factors are identical for ranges at corresponding positions in each of the directional subbands (see Figure 2). The decoding procedure outlined in Table I describes how the mappings from domain to range subtrees may be implemented as operations on individual detail coefficients.

II. GENERALISED LATTICE VECTOR QUANTISATION INTERPRETATION

Considering an image as an element of a vector space, each combination of domain to range assignments generates a manifold within this space, where the coordinates on each manifold consist of the full set of scaling factors and low resolution detail coefficients. The restriction of an image to one of these manifolds constitutes a form of pre-quantisation, with complete quantisation being achieved by quantisation of the coordinates on the manifold. This rep-

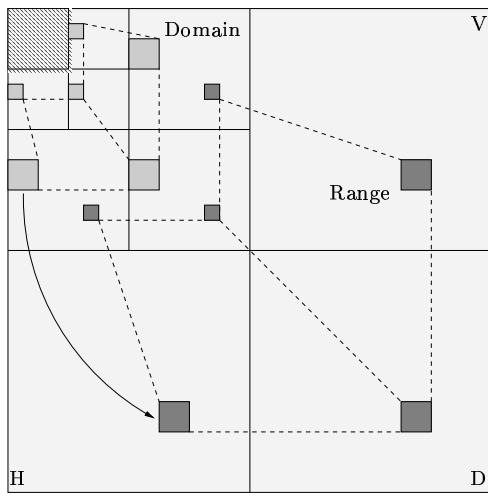


Fig. 2. Domain and range subtrees composed of subtrees rooted at the same resolution in each of the three directional subbands of a non-standard wavelet decomposition.

TABLE I
WAVELET DOMAIN FRACTAL DECODING.

Definitions: (decoding generates an $N \times N$ image)		
i_d	Domain root resolution	$N_d = 2^{i_d}$
$i_r = i_d + 1$	Range root resolution	$N_r = 2^{i_r}$
i_{\max}	Resolution of finest detail	$N = 2^{i_{\max} + 1}$
Fractal representation: $j, k \in \{0, 1, \dots, N_r - 1\}$		
$c_{0;0,0}, d_{i;j,k}$	$0 \leq i \leq i_d$	Low resolution detail
$p_{j,k}, q_{j,k}$	$\in \{0, 1, \dots, N_d - 1\}$	Coordinates of domain mapped to range at coordinates j, k
$s_{j,k}$		Scaling factor applied to domain mapped to range at coordinates j, k
Algorithm:		
For all i from i_r to i_{\max} ($i_r \leq i \leq i_{\max}$):		
$a = \lfloor j2^{i_r - i} \rfloor$	$c = 2^{i - i_r} p_{a,b} + j \bmod (2^{i - i_r})$	
$b = \lfloor k2^{i_r - i} \rfloor$	$d = 2^{i - i_r} q_{a,b} + k \bmod (2^{i - i_r})$	
	$d_{i;j,k} = \frac{1}{2} s_{a,b} d_{i-1;c,d}$	

resentation may be considered as a generalisation of Lattice Vector Quantisation [3], where the reconstruction levels have a far more complex arrangement than a lattice, and are determined by the form of the fractal coding scheme (including, of course, the quantisation applied to the fractal representation) [4, ch. 5]. It is, however, certainly not clear why this should be expected to provide an efficient representation for “natural” images.

Optimal reconstruction vectors for a k -dimensional source \mathbf{X} with pdf $f_{\mathbf{X}}(\mathbf{x})$ have distribution

$$f_{\mathbf{X}}^{\frac{k}{k+2}}(\mathbf{x})$$

in the high resolution case [3, pp. 338-340] and with the MSE distortion measure [3, pg. 471]. Therefore, for large

k , the density of reconstruction points in an optimal codebook should correspond approximately to the probability density of the source. One may therefore assess the accuracy with which a codebook describes a source by comparing corresponding statistics of the source and codebook.

III. STOCHASTIC MODEL

A stochastic signal model exhibiting the form of affine self-similarity assumed by a particular fractal coding scheme may be constructed by considering the elements of the fractal representation (see Table I) as suitably distributed random variables driving the iterative decoding process, just as a random *innovations process* is passed through a linear filter to generate an autoregressive model [5, ch. 2]; this fractal model generates deterministic fractal signals, as opposed to stochastic fractal signals, which may be generated by considerably less complex statistical models, such as multiresolution autoregressive models in which the autoregressive relationship is defined between coefficients at successive resolutions within the tree of detail coefficients in the wavelet transform domain, rather than between successive samples within the domain of the original signal [6]. The density of vectors generated by this model corresponds to the high resolution limit of the related codebook described in the previous section.

Distributions for the low resolution detail coefficients, and scaling factors and domain positions for each range block, were chosen by comparison with the statistics of a set of natural images, since the model statistics are desired to match those of natural images as far as possible. Detail coefficient statistics for natural images were estimated from the wavelet transforms of a set of 18 test images of 512×512 pixels each; the spline wavelet basis [7, table II] used by Davis [2] was used for these calculations.

The lower resolution detail coefficients were found to have an approximately Laplacian distribution, as were the quotients of detail coefficients at consecutive resolutions [4, ch. 6]; a Laplacian distribution was therefore chosen for the initial detail coefficients ($d_{i;j,k}$ for $0 \leq i \leq i_d$) as well as for the scaling factors ($s_{j,k}$). Each domain was allocated equal probability of selection for a particular range ($p_{j,k}$ and $q_{j,k}$) since this distribution is approximately that observed in fractal coding experiments [1, pp. 69-71] (although there is some disagreement in the literature - see reference [4, pp. 46-47, 75-77] for a discussion of this issue).

IV. VARIANCE DECAY

Wavelet domain statistical analysis has been found to be a powerful tool in the analysis of non-stationary signal models such as fractional Brownian motion (fBm) [8]. The variation in detail coefficient variance with increasing resolution is an important statistical property in the wavelet transform domain, corresponding to the power spectral density in many respects - the power spectral density measures signal energy at each frequency, whereas the detail coefficient variance measures signal energy at each resolution.

As a result of the subtree mapping mechanism, each detail coefficient is the product of a subtree scaling factor and a lower resolution detail coefficient, which is itself the product of a scaling factor and a lower resolution detail coefficient, and so on until resolution i_d . Since the same subtree scaling factor may occur multiple times as one traces back this path of influence on a particular detail coefficient (one of the primary differences between this model and multiresolution autoregressive models), the second order statistics of high resolution detail coefficients are dependent² on the higher order statistics [10] of the scaling factors, and detail coefficients at resolution i_d . Since analytic evaluation of the detail coefficient variance for the wavelet domain self-affine models is rather complicated [4, ch. 5], results presented here were generated by Monte Carlo methods. An ensemble of 10000 images of 512×512 pixels was generated by randomly generating the low resolution detail coefficients, scaling factors and domain positions for each range block, decoding each set, as described in Table I, to produce an image.

Measurements of the wavelet transform domain behaviour of the set of test images indicate similar variance decay for the horizontal, vertical and diagonal directional subbands; these values are averaged here for the purpose of comparison with the self-affine model. The log decay with increasing resolution of fBm models provides a good fit to the measured statistics of the test images for resolution 3 and higher. Figure 3 depicts a comparison between this decay and that for the self-affine model with a variance of 4.38×10^5 for the detail coefficients at resolution $i_d = 3$, and a variance of 0.77 for the scaling factors; note the deviation from log decay for the self-affine model. Increasing i_d to 4 or greater (as well as retaining $i_d = 3$ but changing the scaling factor pdf to uniform) reduces the deviation from log decay to the extent that it is barely visible on a graph such as Figure 3. Given suitable restrictions then, the variance decay of the self-affine model may be made compatible with that measured for natural images. All statistics for self-affine models discussed in the following sections are for $i_d = 5$, since this choice provides a good match between model and measured variance decay, and corresponds to range blocks of 8×8 pixels; lower values of i_d correspond to unreasonably large range blocks.

V. INTER- AND INTRASUBBAND CORRELATIONS

Examination of the ensemble of self-affine images indicates that the correlations³ between adjacent detail coefficients in the same subband, averaged over all positions in the subband (stationarity is necessary for this procedure to provide a full description of the correlation structure of the subband), are very close to zero. In comparison, natural images have small but significant correlations between detail coefficients in horizontal and vertical direc-

²Note that the pdfs of scaling factors and initial detail coefficients are not significant if one considers a stochastic model derived from the coding scheme of Rinaldo and Calvagno [9], since independent scaling factors are utilised between each consecutive pair of levels in the coefficient tree.

³Correlation is used here in the sense of normalised covariance.

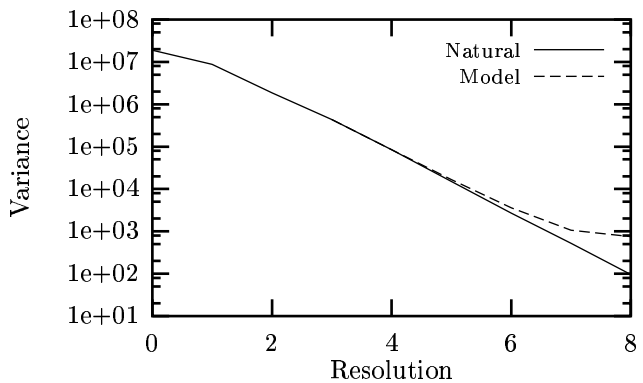


Fig. 3. Measured variance decay for the set of test images, compared with variance decay for the self-affine model with domain roots at resolution 3.

tional subbands, and very small correlations in the diagonal directional subband. The horizontal directional subbands (the subbands containing horizontally oriented edges) have correlations in the region of 0.3 with their neighbours to the left and right, and correlations in the region of -0.2 with their neighbours above and below, while correlations with neighbours in the diagonal directions are considerably smaller. Similar behaviour is observed for the vertical directional subbands, but with reversal of the behaviour for vertical and horizontal neighbours.

Correlations between child and parent coefficients in the self-affine model are very small negative values consistent with a correlation of 0.0. Natural images, in contrast, exhibit a child-parent correlation of approximately 0.1 in the horizontal and vertical directional subbands, and approximately 0.0 in the diagonal directional subbands.

Detail coefficients in both the self-affine model and natural images are therefore, at a rough approximation, uncorrelated. A more accurate view, however, reveals significant differences between the respective correlation structures.

VI. CHILD-PARENT JOINT DISTRIBUTIONS

While smooth basis wavelet transforms, such as that utilised here, are reasonably effective in decorrelating the pixels of natural images, significant dependence remains between parent and child coefficients (while the sign of a child coefficient is approximately uncorrelated with that of its parent, the magnitudes are correlated [11]) - visible in the transformed image by the recurrence of similar patterns of edges in different resolution subbands. Many of the current state-of-the-art image coders, such as the EZW [12] algorithm, were designed to take this dependence into account.

Correlations between the squares of child and parent coefficients are presented in Figure 4 as a measure of the strength of the dependence between coefficient magnitudes. Although increasing with resolution, the correlation is significantly smaller for the self-affine model than in any of the directional subbands of natural images. A simple lin-

ear model of child-parent dependence may be constructed to exhibit the variances, correlations, and correlations between squared coefficients observed for natural images.

ACKNOWLEDGEMENT

The authors wish to thank G. Cox and the anonymous reviewers for valuable comments on the manuscript.

REFERENCES

- [1] Y. Fisher, ed., *Fractal Image Compression: Theory and Application*. Springer-Verlag, 1995.
- [2] G. M. Davis, "A wavelet-based analysis of fractal image compression," *IEEE Transactions on Image Processing*, vol. 7, pp. 141–154, Feb. 1998.
- [3] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*. Kluwer Academic Publishers, 1992.
- [4] B. Wohlberg, *Fractal Image Compression and the Self-Affinity Assumption: A Stochastic Signal Modelling Perspective*. PhD thesis, University of Cape Town, Cape Town, South Africa, Aug. 1996.
- [5] N. S. Jayant and P. Noll, *Digital Coding of Waveforms*. Prentice-Hall, 1984.
- [6] M. Basseville, A. Benveniste, K. C. Chou, S. A. Golden, R. Nikoukhah, and A. S. Willsky, "Modelling and estimation of multiresolution stochastic processes," *IEEE Transactions on Information Theory*, vol. 38, pp. 766–784, Mar. 1992.
- [7] M. Antonini, M. Barlaud, P. Mathieu, and I. Daubechies, "Image coding using wavelet transform," *IEEE Transactions on Image Processing*, vol. 2, pp. 205–220, Apr. 1992.
- [8] P. Flandrin, "Wavelet analysis and synthesis of fractional Brownian motion," *IEEE Transactions on Information Theory*, vol. 38, pp. 910–917, Mar. 1992.
- [9] R. Rinaldo and G. Calvagno, "Image coding by block prediction of multiresolution subimages," *IEEE Transactions on Image Processing*, vol. 4, pp. 909–920, July 1995.
- [10] J. M. Mendel, "Tutorial on higher-order statistics (spectra) in signal processing and system theory: Theoretical results and some applications," *Proceedings of the IEEE*, vol. 79, pp. 278–305, Mar. 1991.
- [11] R. W. Buccigrossi and E. P. Simoncelli, "Progressive wavelet image coding based on a conditional probability model," in *Proceedings ICASSP-97 (IEEE International Conference on Acoustics, Speech and Signal Processing)*, vol. 4, (Munich, Germany), pp. 2597–2600, Apr. 1997.
- [12] J. M. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients," *IEEE Transactions on Signal Processing*, vol. 41, pp. 3445–3462, Dec. 1993.
- [13] B. Wohlberg and G. de Jager, "A review of the fractal image coding literature," To appear in *IEEE Transactions on Image Processing*.

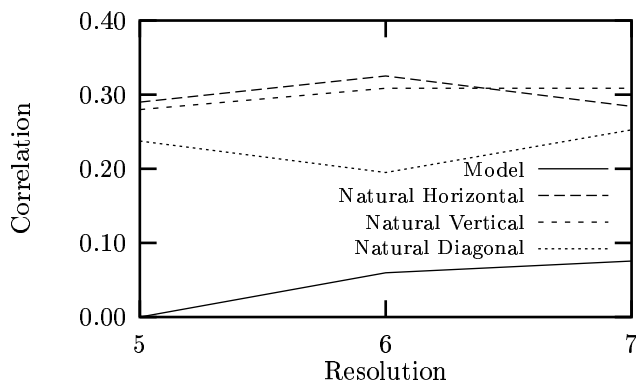



Fig. 4. Correlation between squares of child and parent coefficients for natural images and the self-affine model.

A comparison between the joint probability distributions of child and parent coefficients for natural images and those of the self-affine model is presented in Figure 5, in which substantial differences between the distributions are visible. Noting that the joint distribution of independent Laplacians has diamond-shaped contours, the concavity of the contours for the model reflects a tendency for relatively large parent coefficients to be associated with smaller child coefficients than for an independent distribution (although this tendency is not sufficiently simple to produce a negative correlation between squared coefficients). The convexity of the contours for natural images reflects the tendency for large magnitude parent coefficients to be associated with large magnitude child coefficients, as expected from the correlation between magnitudes described above.


VII. CONCLUSIONS

While appropriate choice of the parameters of the self-affine model allows a good match with the variance decay with increasing resolution observed for natural images, the resulting self-affine model does not accurately represent either the correlation structure of the detail coefficients or the additional dependence between child and parent coefficients. The form of self-affinity considered here therefore does not appear to represent a more accurate characterisation of image statistics than vastly simpler models such as multiresolution autoregressive models.

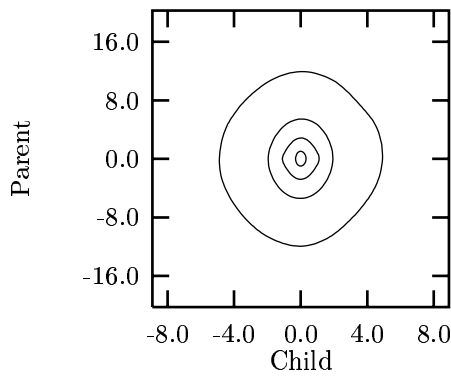
While a number of simplifications have been made in the construction of this model, it does not appear likely that more complex purely self-affine models would provide a significantly improved match with image statistics. This suspicion is supported by the tendency for the most effective fractal coding algorithms not to rely on a purely self-affine representation, but to incorporate aspects of alternative coding strategies [13].



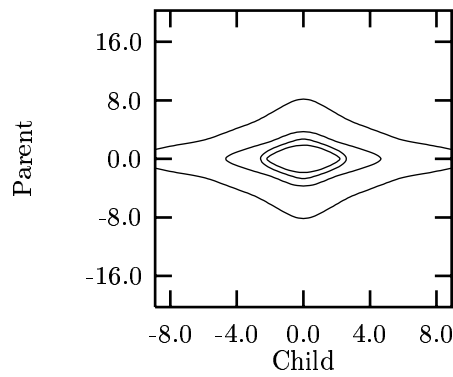
(a) Test images, resolution 5



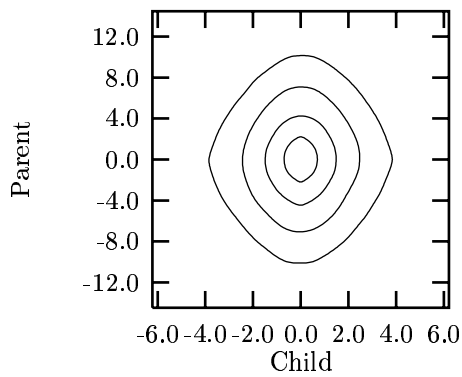
(b) Self-affine images, resolution 5




(c) Test images, resolution 6



(d) Self-affine images, resolution 6



(e) Test images, resolution 7



(f) Self-affine images, resolution 7

Fig. 5. Child-parent joint probability densities for separate ensembles of test images (a distribution for the combined horizontal, vertical, and diagonal directional subbands is given since similar behaviour is observed in each) and images generated by the image model. Distributions are labelled by the resolution of the parent coefficients, and contours are at the same values for each of the three pairs of plots.