A Simple, General Model for the Affine Self-similarity of Images

Simon K. Alexander[1], Edward R. Vrscay[2], and Satoshi Tsurumi[3]

[1] Department of Mathematics, University of Houston, Houston, Texas, USA 77204
[2] Department of Applied Mathematics, Faculty of Mathematics, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1
[3] Department of Information and Computer Engineering, Gunma National College of Technology, 580 Toribamachi, Maebashi, Gunma 371-8530, Japan
skalexander@gmail.com, ervrscay@uwaterloo.ca, tsurumi@ice.gunma-ct.ac.jp

Abstract. A series of extensive numerical experiments indicates that images, in general, possess a considerable degree of affine self-similarity, that is, blocks are well approximated by a number of other blocks – at the same or different scales – when affine greyscale transformations are employed. We introduce a simple model of affine image self-similarity which includes the method of fractal image coding (cross-scale, affine greyscale similarity) and the nonlocal means denoising method (same-scale, translational similarity) as special cases.

1 Introduction

The term "image self-similarity" is subject to a number of possible interpretations which are concerned with how well pixel blocks of an image can, in some way, be approximated by other pixel blocks of the same image. In some applications, such as nonlocal means denoising [3], self-similarity is understood in the strict *translational* sense: Given an image function u and two $n \times n$ pixel blocks R_i and R_j, the two image subblocks $u(R_i)$ and $u(R_j)$ are considered to be "close" only if $u(R_i) \approx u(R_j)$, i.e., the distance $\|u(R_i) - u(R_j)\|$ is small.

From a visual perspective, however, it may be desirable to work with somewhat relaxed requirements. For example, two image subblocks might be considered similar if they are close up to a greyscale shift, i.e., $u(R_i) \approx u(R_j) + \beta$. Consider a picture of a room in which a wall is lit more brightly at one end than at the other. Image blocks from both ends of the wall could be considered to be visually similar. Going a little farther, various "flat" regions of an image, e.g., a wall, a clear sky, a table, could be classed as visually similar.

A further relaxation is afforded by allowing affine greyscale transformations, i.e., $u(R_i) \approx \alpha u(R_j) + \beta$. For example, in "structured vector quantization using linear transforms" [7], image blocks are approximated by affinely transformed codebook blocks. In fractal image coding, one approximates blocks of an image across scales: $u(R_i) \approx \alpha u(D_j) + \beta$, where D_j is larger than R_i.

In this paper, we introduce a simple model of *local affine image self-similarity* that accomodates all of the above examples as special cases. The original motivation arises from our work in fractal image coding, in particular some recent work

A. Campilho and M. Kamel (Eds.): ICIAR 2008, LNCS 5112, pp. 192–203, 2008.

demonstrating its effectiveness in image denoising [1,10,11]. We have also been inspired by the increasing interest in nonlocal methods of image processing that exploit self-similarity, for example, restoration [19], denoising [3,4] and zooming [6,9] – see also [5].

Our investigation is centered around a series of extensive numerical experiments that examine the distributions of errors in approximating image blocks $u(R_i)$ by affine greyscale transformations of other image blocks $u(D_j)$. As noise of increasing variance is added to an image, its domain-range error distribution will be shifted outward. In the limit of zero-signal-to-noise ratio, the distribution of the pure noise image can be characterized analytically. These results provide a partial answer to the following question posed by D.L. Ruderman [18], "In which ways do natural images differ from random images?"

Images with error distributions that are more concentrated near zero error may be viewed as possessing greater degrees of self-similarity. This suggests that relative degrees of self-similarity can be characterized quantitatively in terms of the means and variances of the error distributions. We present computational results for some standard test images. Our results provide some explanation of why self-similar-based methods, including fractal image coding, work so well, approximating or denoising images quite effectively.

Finally, we show that the error distribution of an image is generally similar to the distribution of block variances of the image. Since flatter blocks are more easily approximated, one could argue that "image self-similarity" could be replaced by the term "image approximability."

2 A Simple Class of Models for Image Self-similarity

An image I will be represented by an image function $u : X \to R_g$, where $R_g \subset \mathbf{R}$ denotes the *greyscale range*. In the computations presented below, we work with normalized images, i.e., $R_g = [0, 1]$, converting them to 8 bit-per-pixel images for display. The support X of an image function u is assumed to be an $n_1 \times n_2$-pixel array. The components of our model are as follows:

1. A set $\mathcal{R}$ of $n \times n$-pixel *range* subblocks R_i, $1 \le i \le N_R$ such that (i) $R_i \cap R_j = 0$ if $i \ne j$ and (ii) $X = \cup_i R_i$. In other words, $\mathcal{R}$ forms a partition of X. We let $u(R_i)$ denote the portion of u that is supported on R_i.
2. A set $\mathcal{D}$ of $m \times m$-pixel *domain* subblocks D_j, where $m \ge n$. The set of blocks $\mathcal{D}$ should cover X, i.e., $\cup_j D_j = X$ but they need not be nonoverlapping.
3. The geometric transformations $w_{ij}^{(k)}$ that map a domain block D_j to range block R_i. For simplicity, we consider only *affine* transformations. Since both blocks are square, there are 8 possible mappings (four rotations and four inversions about the center) which are accomodated in the index $1 \le k \le 8$. In the case that $m > n$, i.e., D_j is larger than R_i, it is also assumed that the *contractive* map w_{ij} includes an appropriate pixel decimation operation.
4. Affine greyscale maps $\phi : R_g \to R_g$ having the form $\phi(t) = \alpha t + \beta$.

Given an image function u, we examine how well or poorly the subimages $u(R_i)$ are approximated by subimages $u(D_j)$, to be written symbolically as

$$u(R_i) \approx \phi_{ij}(u(D_j)) = \alpha_{ij} u(D_j) + \beta_{ij}, \quad 1 \le i \le N_R, \;\; 1 \le j \le N_D, \qquad (1)$$

with the understanding that the relation applies at the single pixel level. (Technically, the above should be written as $u(R_i) \approx \phi_i(u(w_{ij}^{-1}(R_j))) \cdots$. Note that we have also omitted the k superscripts for convenience.)

We emphasize that the model formulated above has been made as simple as possible. As such, we address some potential concerns briefly below:

1. The use of *square, nonoverlapping* blocks: This was an effort to standardize the method, with low computational cost. Generally, the same behaviour is observed for larger numbers of overlapping blocks.
2. The use of range blocks of the *same size*: Generally, the smaller a block, the easier it is to approximate it. We are attempting to keep all regions of an image "on the same playing field." One may certainly wish to examine the self-similarity statistics over several scales, i.e., sizes of range blocks.
3. The use of *affine* greyscale maps $\phi(t) = \alpha t + \beta$. Such a family of maps is very simple in form yet, with two parameters, sufficiently flexible. Of course, better approximations would be accomplished with higher-degree polynomials but we believe that such similarities would be artificial.

In an effort to characterize how well images may be approximated with this model, we consider the distribution of errors Δ_{ij} associated with Eq. (1), i.e.,

$$\Delta_{ij} = \min_{\alpha,\beta \in \Pi} \|u(R_i) - \alpha u(D_j) - \beta\|, \quad 1 \le i \le N_R, \;\; 1 \le j \le N_D. \qquad (2)$$

Here, $\|\cdot\|$ denotes the $L^2(X)$ norm. In all calculations reported in this paper, the L^2-distance between two $n \times n$ image subblocks $u(R_i)$ and $v(R_i)$ will be the root-mean-square (RMS) distance. $\Pi \subset \mathbf{R}^2$ denotes the feasible (α, β) parameter space, restricted so that $\phi : R_g \to R_g$.

We consider four particular cases of this self-similarity model:

1. **Purely translational:** Domain and range blocks have the same size, i.e., $m = n$. As such, the w_{ij} are translations and $\alpha_{ij} = 1$, $\beta_{ij} = 0$. The approximation error is simply $\Delta_{ij} = \|u(R_i) - u(D_j)\|$.
2. **Translational + greyscale shift:** The w_{ij} are again translations. We set $\alpha_{ij} = 1$ and optimize over β. The approximation error is $\Delta_{ij} = |\beta_{ij}| = [\bar{u}(R_i) - \bar{u}(D_j)]$, where the bars denote mean values of the subblocks.
3. **Affine, same-scale:** The w_{ij} are translations and we optimize over α, β.
4. **Affine, two-scale:** The w_{ij} are affine spatial contractions (which involve decimations in pixel space). We optimize over α, β.

3 Cases 1,2 and 3: Same-Scale Self-similarity

Here, the domain and range blocks have the same size. We naturally expect that for a given domain-range pairing (D_j, R_i), the approximation errors of Eq. (2) for Cases 1, 2 and 3 will behave as follows:

$$0 \leq \Delta_{ij}^{(Case\ 3)} \leq \Delta_{ij}^{(Case\ 2)} \leq \Delta_{ij}^{(Case\ 1)}, \tag{3}$$

since one optimizes over more parameters as we move from Case 1 (no parameters) to Case 2 (one parameter) to Case 3 (two parameters). In the numerical experiments reported below, the domain and range blocks were taken from the same set of nonoverlapping 8×8-pixel blocks, i.e., $D_i = R_i$. Fig. 1 summarizes the results of calculations on the normalized *Lena* and *Mandrill* images.

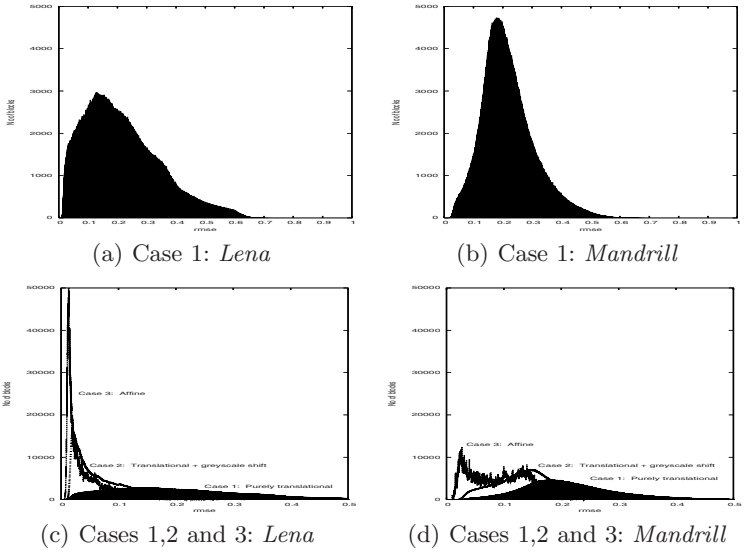

(a) Case 1: *Lena* (b) Case 1: *Mandrill*

(c) Cases 1,2 and 3: *Lena* (d) Cases 1,2 and 3: *Mandrill*

Fig. 1. Same-scale RMS self-similarity error distributions for normalized *Lena* and *Mandrill* images In all cases, 8×8-pixel blocks R_i and D_j were used

The top row shows histogram distributions of the approximation errors for Case 1, over the interval $[0, 1]$. These errors are simply the RMSE distances, $\Delta_{ij}^{(Case\ 1)} = \|u(R_i) - u(R_j)\|$. At first glance, one may well surmise that these images are quite similar translationally: Both distributions exhibit significant peaking at around $\Delta = 0.15$, with the *Mandrill* image being more pronounced.

The bottom row of Fig. 1, however, shows that enormous reductions in the approximation error are achieved when one employs greyscale maps, even for Case 2, where only a greyscale shift β is used – note that the distributions are plotted over the subinterval $[0, 0.5]$. The Case 1 Δ-error distributions (shaded) are included in these plots for comparison.

In Fig. 2 are plotted the histogram distributions of the standard deviations $\sigma(u(R_i))$ of the 8×8 range blocks. There is a noteworthy similarity between these distributions and the Case 2 distributions of Fig. 1 which can be explained as follows. The standard deviation of an image block $\sigma(u(R_i))$, is the RMSE error in approximating $u(R_i)$ by its mean value, $\bar{u}(R_i)$. This is equivalent to setting

the greyscale parameter $\alpha = 0$ and optimizing over β in Eq. (1). Removing the condition $\alpha = 0$ will generally produce better approximations, i.e.,

$$0 \le \Delta_{ij}^{(Case\ 2)} \le \sigma(u(R_i)). \tag{4}$$

As such, the Case 2 Δ-error distributions will be shifted perturbations of the block variance distributions.

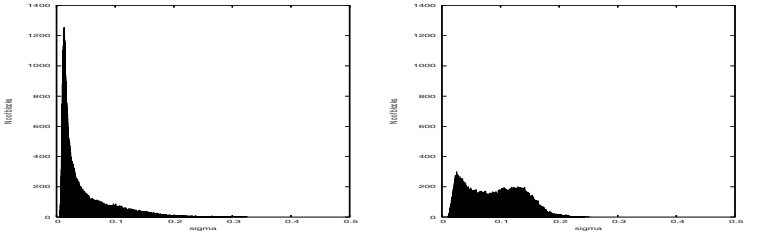

(a) Lena (b) Mandrill

Fig. 2. Distributions of $\sigma(u(R_i))$ of 8×8-pixel blocks for normalized *Lena* and *Mandrill* images, over the interval $[0, 0.5]$. Note the similarity to Case 2 distributions of Fig. 1.

Although there does not seem to be much difference between *Lena* and *Mandrill* in terms of translational similarity (Case 1), a significant difference is produced when a greyscale shift β is used (Case 2). This can be explained as follows. Fig. 2 shows that the *Lena* image contains a significantly higher proportion of "flatter" image subblocks, i.e., blocks of low variance, than the *Mandrill* image. From Eq. (4), the Case 2 Δ-error distribution for *Lena* will be more concentrated near zero. Further improvement is expected with Case 3, cf. Eq. (3).

The above discussion, in particular Eq. (4), suggests that the distribution of block variances is the most important factor in how well subblocks of an image I may be approximated by other subblocks, i.e., its degree of "self-similarity." An extreme example is the constant image $u = C$. Here, the Δ-error distribution consists of a single peak at $\Delta = 0$. We return to this idea in a later section.

Application to "nonlocal means denoising". As is well known, a standard technique for the reduction of additive white noise is to average over multiple samples. This is the basis of the very effective "nonlocal-means denoising algorithm" [3], where the multiple samples are provided by the image itself. Very briefly, each pixel $u(i)$ of a noisy image is replaced by a convex combination of other pixel values $u(j)$ from the image. The weights λ_{ij} of this averaging procedure depend upon the similarity between neighbourhoods N_i and N_j centered about pixels i and j, respectively. Neighbourhoods N_k that do not approximate N_i very well, i.e., with high L^2 error $\|N_i - N_k\|$, are assigned low weights. In essence, the nonlocal-means algorithm relies on the translational self-similarity of an image, i.e., Case 1.

It is remarkable that the nonlocal-means denoising method works so well. Because of the translational symmetry requirement, only a few blocks generally contribute significantly to the denoising of a given pixel. In some applications, it would not seem unreasonable to relax this restriction and allow constant greyscale shifts (Case 2), thereby increasing significantly the number of blocks that could contribute to the denoising. This slight relaxation of the method is observed to improve denoising. Moreover, the computational cost is minimal since the optimal greyscale shifts β are easily computed.

4 Case 4: Two-Scale, Affine Self-similarity

This is the essence of fractal image coding [8,17]. Given a "target" image u, each subblock $u(R_i)$ is approximated by a geometrically-contracted, affine greyscale-modified copy of a larger subblock $u(D_j)$. The range-domain assignments $(i, j(i))$ and associated optimal greyscale parameters (α_i, β_i) comprise the *fractal code* of u that defines a *fractal transform* operator T. Eq. (1) then becomes

$$u(x) \approx (Tu)(x) = \alpha_i u(w_{i,j(i)}^{-1}(x)) + \beta_i, \quad x \in R_i, \quad 1 \le i \le N_R. \tag{5}$$

Under appropriate conditions involving the α_i and the contraction factors c_i of the spatial maps $w_{i,j(i)}$, T is contractive in $L^2(X)$. Then, from Banach's Fixed Point Theorem, there exists a unique fixed point function $\bar{u} = T\bar{u}$. Furthermore, $\bar{u}$ may be generated by the iteration procedure $u_{n+1} = Tu_n$, where u_0 is any "seed" image: $u_n \to \bar{u}$ as $n \to \infty$. (The convergence is geometric.) By construction, $\bar{u}$ is an approximation to the target image u.

The fractal transform operator T in Eq. (5) is a *nonlocal* operator since blocks of an image function u are approximated with modified copies of blocks from elsewhere in the image. For this reason, fractal coding has often been referred to as "self-vector quantization". The connection between fractal coding and vector quantization was realized many years ago, e.g. [16,15,14,13]. Perhaps it is more appropriate to consider fractal coding as a "self-structured VQ using linear transforms," cf. [7].

The mathematical basis for this method of approximation is provided by the so-called *Collage Theorem* [2], a simple consequence of Banach's Theorem:

$$\|u - \bar{u}\| \le \frac{1}{1 - c_T} \|u - Tu\|, \tag{6}$$

where c_T is the contraction factor of T. Given a set of range blocks $\mathcal{R}$ and domain blocks $\mathcal{D}$, one tries to make the approximation error $\|u - \bar{u}\|$ small by minimizing the *collage error* $\|u - Tu\|$. From Eq. (5), this is done as follows: For each range block R_i, we search the domain pool $\mathcal{D}$ for the block $D_{i,j(i)}$ that yields the lowest collage error Δ_{ij} in Eq. (2).

In Fig. 3 is presented the fixed point approximation $\bar{u}$ to the standard 512×512-pixel *Lena* image (8 bits per pixel) obtained using 8×8-pixel range blocks. The domain pool for each range block was the set of $32^2 = 1024$ nonoverlappping

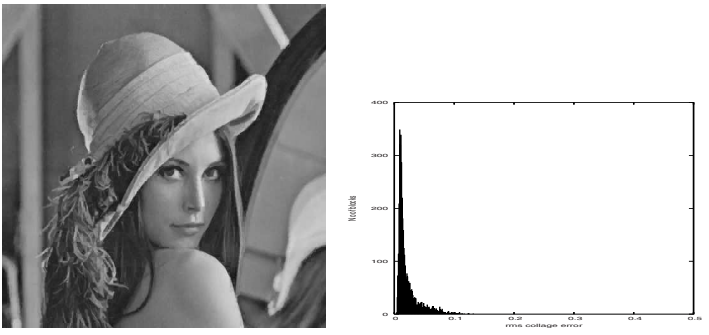

(a) $\bar{u}$, PSNR = 30.37 dB (b) $\Delta_{i,j(i)}$

Fig. 3. (a) Fixed point $\bar{u}$ approximation to *Lena* image, as discussed in text. (b) Distribution of RMS collage errors $\Delta_{i,j(i)}$ (normalized image) over $[0, 0.5]$ for the $64^2 = 4096$ domain-range pairs defining the fractal transform T.

16×16-pixel blocks. (A better approxmation could be obtained with the use of a larger domain pool, but at the expense of computational search time.) Also presented in Fig. 3 is the distribution of all RMS collage errors $\Delta_{i,j(i)}$ between between range and selected domain blocks used to define the fractal transform operator T. The significant peak of collage errors near zero error indicates that a large fraction of range blocks is well approximated by this procedure.

We now examine how well/poorly *all range blocks* are approximated by *all possible domain blocks*. The histogram distribution of all such possible collage errors Δ_{ij} for the *Lena* image is presented in Fig. 4(a). This distribution also demon-

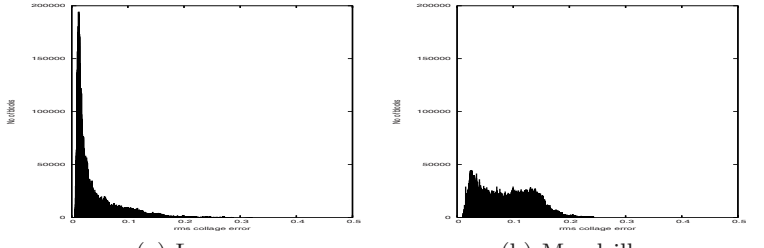

(a) Lena (b) Mandrill

Fig. 4. RMS collage error distributions over $[0, 0.5]$ for normalized *Lena* and *Mandrill* images. 8×8-pixel range blocks.

strates a significant peak near zero, indicating that a majority of domain-range pairings yield low error. There is also a great similarity between this distribution and that of the same-scale case (Case 3) in Fig. 1(c). In both same-scale and cross-scale cases, a given range block is generally well-approximated by a number of domain blocks.

Fig. 4 also shows the distribution of collage errors for the *Mandrill* image. This distribution is more diffuse than that of the *Lena* image, indicating that range blocks of *Mandrill* are not as well approximated. This is consistent with our observations in the same-scale case, cf. Fig. 1(d). Note also the similarity between these two distributions and the distributions of σ-values for 8×8-pixel blocks for the same images in Fig. 2. Recall that such a similarity is to be expected: The standard deviation σ of an image range block $u(R_i)$ is the error in approximating it with the constant value $\bar{u}(R_i)$, which corresponds to setting the greyscale parameter $\alpha = 0$ and optimizing over β.

Historically, most fractal image coding research focussed on its compression capabilities – obtaining acceptable accuracy with the smallest possible domain pool. Understandably, these investigations would rarely venture beyond observing what the "optimal" domain blocks would provide. Our study is not concerned with the rate-distortion properties of fractal coding, but rather with the degree of self-similarity in images, as reflected by the *redundancy* of good domain-range matchings. That being said, the former is certainly influenced by the latter.

5 Effects of Noise

One expects that the presence of noise in an image will decrease the ability of its subblocks to be approximated by other subblocks. Because of our primary interest in fractal coding, results are presented below for the two-scale case (Case 4). Similar behaviour is exhibited for single-scale similarity (Cases 1-3). In what follows we let I_0 denote a noiseless test image, to which zero-mean Gaussian noise $\mathcal{N}(0, \sigma^2)$ is added to produce a noisy image $I(\sigma)$.

Some simple experiments show that as noise $\mathcal{N}(0, \sigma^2)$ of increasing variance σ^2 added to an image, the peak of the distribution of collage errors Δ_{ij} moves away from zero. Moreover, the Δ-error distribution becomes more diffuse. These features are demonstrated in Fig. 5 for the two cases of the normalized *Lena* image plus noise with σ values of 0.1 and 0.3 For comparison purposes, we have also plotted the Δ-error distributions for "pure noise images" of the form $n(\sigma) = 0.5 + \mathcal{N}(0, \sigma^2)$. (For a given σ value, $n(\sigma)$ may be considered as a kind of "zero signal-to-noise limit" of a noisy *Lena* image.) Note that the Δ-error distributions of $n(\sigma)$ are situated roughly at σ. This can be shown analytically: When we approximate noise blocks with other noise blocks, the optimal values of the greyscale parameters (at least for infinite block size) are $\alpha = 0$ and $\beta = \sigma$. Note also that the width of the distribution also increases with σ.

It is noteworthy that the Δ-error distributions for the noisy *Lena* images also peak at the σ values of the added noise, even for the relatively mild case $\sigma = 0.1$. In fact, the Δ-error distribution for *Lena* + $\mathcal{N}(0, 0.3^2)$ is virtually identical to that of the pure noise case. Note as well that the distribution for *Lena* + $\mathcal{N}(0, 0.3^2)$ is not as concentrated about the peak as the distribution for the pure noise image $n(0.1)$. At such a low σ-value, the image $n(\sigma)$ is roughly constant, or at least more statistically constant than the *Lena* image. In this case, $n(\sigma)$ could therefore be viewed as more self-similar than the *Lena* image.

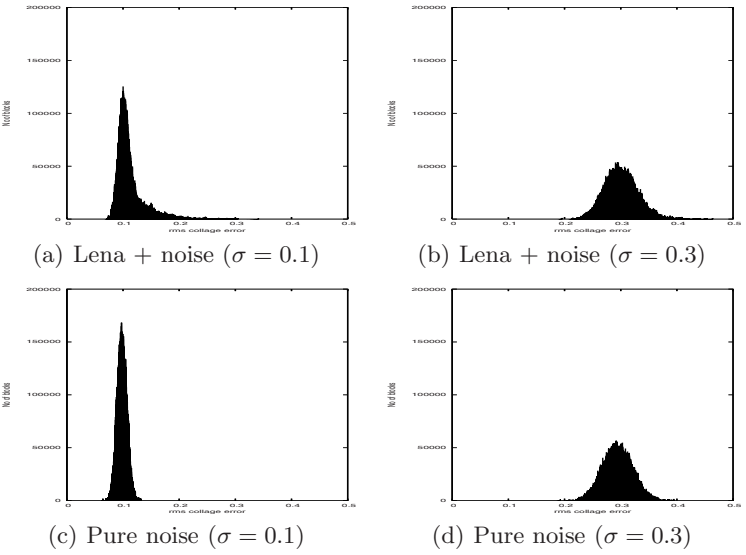

(a) Lena + noise ($\sigma = 0.1$) (b) Lena + noise ($\sigma = 0.3$)

(c) Pure noise ($\sigma = 0.1$) (d) Pure noise ($\sigma = 0.3$)

Fig. 5. Distributions of collage errors Δ_{ij} over $[0, 0.5]$ for two cases of *Lena* image + zero-mean, Gaussian noise, along with distributions of pure noise images for comparison. 8×8-pixel range blocks, 16×16-pixel range blocks.

The coincidence of the peaks of the noisy *Lena* images $I(\sigma)$, e.g. *Lena*, and their pure noise counterparts $n(\sigma)$ actually illustrates a simple and standard method of estimating the variance σ^2 of additive noise in a noisy image: one simply constructs a histogram of the local block variances and notes the location of the peak [12].

Fractal image denoising. As with any lossy compression method, simple fractal coding of a noisy image $I(\sigma)$ produces some denoising [10,11] . There are two principal reasons: (i) the affine greyscale fitting between domain and range blocks causes some variance reduction in the noise, and (ii) the spatial contraction/pixel decimation involved in mapping domain blocks to range blocks provides further variance reduction. Additional denoising can be obtained by using estimates of the noise variance to estimate the fractal code of the noiseless image [10].

The fact that each range block is well approximated by a number of domain blocks can be exploited to perform denoising by using multiple copies [1], a cross-scale analog of the nonlocal means denoising method. Space limitations preclude a more detailed description of this *multi-parent fractal transform method*.

6 Using Δ-Error Distributions to Assign Relative Self-similarity

Let us now return to the idea of using collage error distributions to characterize the degree of self-similarity in images. We have computed the Δ-error

distributions of a large number of test images, to find that they lie across – and even beyond – the spectrum spanned by the *Lena* and *Mandrill* images. The Δ-error distributions of a few standard (512×512-pixel) images are presented in Fig. 6. Note that the distributions of *San Francisco*, *Boat*, *Peppers* and *Barbara* strongly resemble that of *Lena*, whereas the *Zelda* distribution leans much more toward *Mandrill*, with that of *Goldhill* not far behind.

The means and standard deviations of the collage error distributions for the test images are listed in Table 1. The entries have been arranged in a kind of "decreasing self-similarity" based upon increasing mean and, to some extent, increasing width. Estimates of the (natural logarithm) entropies of these distributions have also been presented in this table (third column) – note that with the exception of *San Francisco*, they increase as we proceed down the table.

In the final two columns of this table we present the estimates of the means and standard deviations of the distributions of standard deviations for these

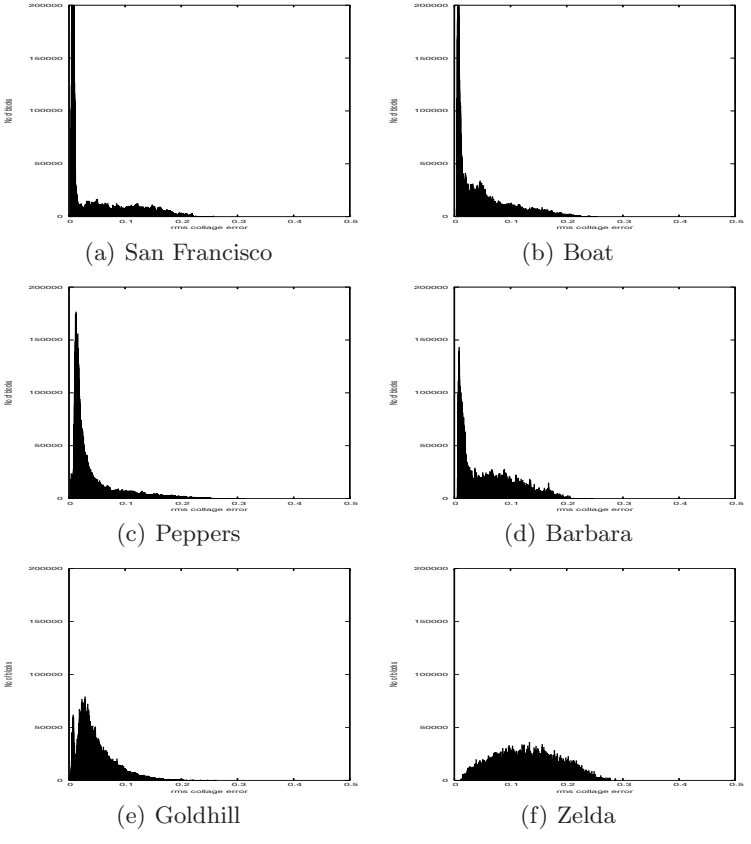

Fig. 6. RMS collage error distributions over $[0, 0.5]$ for some other (normalized) test images: 8×8-pixel range blocks. All possible domain-range pairs were considered along with eight spatial mappings.

Table 1. Columns 1-3: Means, standard deviations, and entropies of collage error distributions for test images examined in this paper. Columns 4 and 5: Means and standard deviations of subblock σ-distributions of these images, to show their agreement with Columns 1 and 2, respectively.

Image	Collage errors			Range block stddevs	
	mean	stddev	entropy	mean	stddev
Lena	0.043	0.044	2.26	0.046	0.046
San Francisco	0.046	0.057	2.01	0.048	0.059
Peppers	0.047	0.050	2.32	0.049	0.052
Goldhill	0.049	0.034	2.46	0.052	0.036
Boat	0.052	0.052	2.58	0.055	0.055
Barbara	0.060	0.049	2.69	0.064	0.051
Mandrill	0.089	0.048	2.85	0.089	0.048
Zelda	0.126	0.055	3.09	0.141	0.054

images, cf. Fig. 2, to show their excellent agreement with those of the collage error distributions.

Finally, there may still be a concern that the images examined above do not form a suitably broad sampling of "natural images." For this reason, the experiments have been repeated on a much larger set of natural images, 700 images from 21 datasets in total taken from the University of Washington 'Groundtruth Database'. Our findings are qualitatively similar to those reported above.

7 "Self-similarity" vs "Approximability"

At the end of Section 2, we observed that the degree of self-similarity of an image is determined primarily by the distribution of its *block variances*. To pursue this idea further, we have examined numerically how the $n \times n$-pixel blocks of an image A are approximated, under affine greyscale transformations, by $n \times n$-pixel blocks of *another image* B. We find, in general, that the resulting error distribution is virtually identical to that of approximating blocks of A with other blocks of A. For example, in the cases $A = Lena$ and $B = Mandrill$ and *vice-versa*, we obtain error distributions virtually identical to the Case 3 distributions of Fig. 1(c) and 1(d), respectively. (We omit the actual plots because of space limitations.) This phenomenon is also observed for (cross-scale) fractal image coding: In the cases $A = Lena$ and $B = Mandrill$ and *vice-versa*, we obtain collage error distributions that are virtually identical to those of Fig 4.

These observations indicate that the *source* of domain blocks for an image is not as important as the *ability to approximate* the range blocks of the image. We therefore conclude that *the degree of self-similarity of an image is a consequence of how well its range blocks can be approximated.* As we have seen, the latter can be decided on the basis of the variance distribution of the range blocks.

Acknowledgements

We gratefully acknowledge the generous support of this research by the Natural Sciences and Engineering Research Council of Canada (NSERC) in the forms of a Discovery Grant (ERV) and a Postgraduate Scholarship and Postdoctoral Fellowship (SKA). ST would also like to acknowledge the support of the Ministry of Education, Culture, Sports, Science and Technology in Japan which made possible his visit to Waterloo as a Fellow for Research Abroad (2002-2003).

References

1. Alexander, S.K.: Multiscale Methods in Image Modelling and Image Processing, Ph.D. Thesis, Dept. of Applied Mathematics, University of Waterloo (2005)
2. Barnsley, M.F.: Fractals Everywhere. Academic Press, New York (1988)
3. Buades, A., Coll, B., Morel, J.M.: A review of image denoising algorithms, with a new one. Multiscale Modelling and Simulation 4, 490–530 (2005)
4. Dabov, K., Foi, A., Katkovnik, V., Egiazarian, K.: Image denoising by sparse 3-D transform-domain collaborative filtering. IEEE Trans. Image Proc. 16, 2080–2095 (2007)
5. Ebrahimi, M., Vrscay, E.R.: Solving the Inverse Problem of Image Zooming Using "Self-Examples". In: Kamel, M., Campilho, A. (eds.) ICIAR 2007. LNCS, vol. 4633, pp. 117–130. Springer, Heidelberg (2007)
6. Elad, M., Datsenko, D.: Example-based regularization deployed to super-resolution reconstruction of a single image. The Computer Journal 50, 1–16 (2007)
7. Etemoglu, C., Cuperman, V.: Structured vector quantization using linear transforms. IEEE Trans. Sig. Proc. 51, 1625–1631 (2003)
8. Fisher, Y. (ed.): Fractal Image Compression: Theory and Application. Springer, New York (1995)
9. Freeman, W.T., Jones, T.R., Pasztor, E.C.: Example-based super-resolution. IEEE Comp. Graphics Appl. 22, 56–65 (2002)
10. Ghazel, M., Freeman, G., Vrscay, E.R.: Fractal image denoising. IEEE Trans. Image Proc. 12, 1560–1578 (2003)
11. Ghazel, M., Freeman, G., Vrscay, E.R.: Fractal-wavelet image denoising. IEEE Trans. Image Proc. 15, 2669–2675 (preprint, 2006)
12. Gonzalez, R.C., Woods, R.E.: Digital Image Processing. Prentice-Hall, New Jersey (2002)
13. Hamzaoui, R.: Encoding and decoding complexity reduction and VQ aspects of fractal image compression, Ph.D. Thesis, University of Freiburg (1998)
14. Hamzaoui, R., Müller, M., Saupe, D.: VQ-enhanced fractal image compression. In: ICIP 1996. IEEE, Los Alamitos (1996)
15. Hamzaoui, R., Saupe, D.: Combining fractal image compression and vector quantization. IEEE Trans. Image Proc. 9, 197–207 (2000)
16. Lepsoy, S., Carlini, P., Oien, G.: On fractal compression and vector quantization. In: Fisher, Y. (ed.) Fractal Image Encoding and Analysis. NATO ASI Series F, vol. 159. Springer, Heidelberg (1998)
17. Lu, N.: Fractal Imaging. Academic Press, New York (1997)
18. Ruderman, D.L.: The statistics of natural images. Network: Computation in Neural Systems 5, 517–548 (1994)
19. Zhang, D., Wang, Z.: Image information restoration based on long-range correlation. IEEE Trans. Cir. Syst. Video Tech. 12, 331–341 (2002)