Self-similarity of Images in the Wavelet Domain in Terms of ℓ^2 and Structural Similarity (SSIM)

D. Glew and Edward R. Vrscay

Department of Applied Mathematics, Faculty of Mathematics,
University of Waterloo, Waterloo, Ontario, Canada N2L 3G1
{dtglew,ervrscay}@uwaterloo.ca

Abstract. Images exhibit a high degree of affine self-similarity with respect to the L^2 distance. That is, image subblocks are generally well-approximated in L^2 by a number of other (affine greyscale modified) image subblocks. This is due, at least in part, to the large number of flatter blocks that comprise such images. These blocks are more easily approximated in the L^2 sense, especially when affine greyscale transformations are employed. In this paper, we show that wavelet coefficient quadtrees also demonstrate a high degree of self-similarity under various affine transformations in terms of the ℓ^2 distance. We also show that the approximability of a wavelet coefficient quadtree is determined by the lowness of its energy (ℓ^2 norm).

In terms of the structural similarity (SSIM) index, however, the degree of self-similarity of natural images in the pixel domain is not as high as in the L^2 case. In essence, the greater approximability of flat blocks with respect to L^2 distance is taken into consideration by the SSIM measure. We derive a new form for the SSIM index in terms of wavelet quadtrees and show that wavelet quadtrees are also not as self-similar with respect to SSIM. In an analgous way, the greater approximability of low-energy quadtrees is taken into consideration by the wavelet-based SSIM measure.

1 Introduction

In [2], a simple model of affine self-similarity of images was introduced, based upon the observation that pixel subblocks of natural images are generally well approximated in L^2 distance by other pixel subblocks. This simple model included various degrees of affine greyscale transformations as well as spatial transformations, thereby including a number of nonlocal image processing methods as special cases, in particular, nonlocal-means denoising [5] and fractal image coding [8,11].

In this paper, we investigate the self-similarity of images in the wavelet domain, observing how well wavelet coefficient quadtrees are approximated by other quadtrees at the same or different levels. Because our formalism should, in principle, be applicable to quadtrees of infinte length, only scaling transformations of the form $\mathbf{y} = \alpha\mathbf{x}$ are allowed in order to preserve ℓ^2 summability. Our observations and conclusions will follow those of [2] and [3] in the pixel domain

A. Campilho and M. Kamel (Eds.): ICIAR 2012, Part I, LNCS 7324, pp. 131–140, 2012.

rather closely, i.e., that in terms of ℓ^2 distance, wavelet coefficient quadtrees can demonstrate a significant degree of self-similarity. However, when examined in terms of a wavelet-based SSIM measure, which we define for the first time in this paper, the degree of self-similarity is diminished.

In order to understand these results, it is necessary to review very briefly the main aspects and results of the "pixel-based" approach in [2] and [3]. First of all, let $\Delta_{ij}^{(k)}$ denote the L^2-error in approximating an image "range" block $u(R_i)$ by a "domain" block $u(D_j)$, where R_i and D_j have, unless otherwise indicated, the same size (typically 8×8 pixel blocks). The superscript k indicates the type, or "Case", of greyscale transformation employed in the approximation,

- **Case 1: Purely translational**, where $u(R_i) \approx u(D_j)$. This is the basis of nonlocal means denoising.
- **Case 2: Translational + greyscale shift**, where $u(R_i) \approx u(D_j) + \beta_{ij}$.
- **Case 3: Affine transformation**, where $u(R_i) \approx \alpha_{ij} u(D_j) + \beta_{ij}$.
- **Case 4: Affine transformation,**, as in Case 3, but **cross-scale**, i.e., the domain block D_j is larger. This is the basis of fractal image coding.

The Case 3 Δ-error distributions of some images such as *Lena*, *San Francisco* and *Boat* were observed to be quite concentrated near zero error. Other distributions, e.g., from *Barbara* and *Mandrill*, were more diffuse. The Case 3 Δ-error distribution of an image, however, is observed to be quite similar in structure to the distribution of the standard deviations of subblocks in the image. This is not surprising when it is recalled that the standard deviation of a block $u(R_i)$ is the L^2-error associated with the best constant approximation, namely its mean value $\overline{u(R_i)}$. The distributions of α coefficients in the Case 3 affine transformation are strongly peaked at $\alpha = 0$ for all images, indicating that the this approximation will usually be a very slight improvement over the approximation of a subblock by its mean.

From these observations, the question was raised whether the "self-similarity" of an image is actually due to the *approximability* of its subblocks which, in the L^2 (least-squares) sense, is determined by their "flatness," i.e., lowness of standard deviation σ or variance σ^2. Images with a higher proportion of flatter, i.e., low-variance, blocks will exhibit a greater degree of peaking near zero, particularly for the case $k = 3$.

In [3] the self-similarity of images was investigated in terms of the Structural Similarity (SSIM) index [13,14]. It was found that SSIM removes the "unfair advantage" of low-variance blocks. For images with high proportions of low-variance blocks, the peaking of the distributions of SSIM indices between range blocks is not as dramatic as in the L^2 case. Nevertheless, the degree of peaking of SSIM distributions does vary between images but in a less pronounced manner. In terms of SSIM, the *Lena* image is still more self-similar than *Mandrill*, although their SSIM distributions do not differ as much as their L^2 counterparts.

Here, we briefly mention how the "unfair advantage" of low-variance blocks is removed by SSIM since the same idea will apply to wavelet quadtree self-similarity. If we let $\mathbf{x}, \mathbf{y} \in \mathbb{R}_+^N$ denote two nonnegative image subblocks, then

their SSIM index $S(\mathbf{x}, \mathbf{y}) = 1$ if and only if $\mathbf{x} = \mathbf{y}$. This suggests that the function

$$T(\mathbf{x}, \mathbf{y}) = 1 - S(\mathbf{x}, \mathbf{y}), \qquad \mathbf{x}, \mathbf{y} \in \mathbb{R}_+^N, \tag{1}$$

could be a measure of the distance between $\mathbf{x}$ and $\mathbf{y}$ since $\mathbf{x} = \mathbf{y}$ implies that $T(\mathbf{x}, \mathbf{y}) = 0$. In fact, $T(\mathbf{x}, \mathbf{y})$ is related to the L^2 distance $\|\mathbf{x} - \mathbf{y}\|_2$ as follows [3].

First of all, if $\mathbf{x}$ represents an image range block $u(R_i)$ and $\mathbf{y}$ its optimal affine (Case 3) optimization then $\bar{\mathbf{x}} = \bar{\mathbf{y}}$. In this case,

$$T(\mathbf{x}, \mathbf{y}) = \frac{1}{N-1} \frac{\|\mathbf{x} - \mathbf{y}\|_2^2}{s_{\mathbf{x}}^2 + s_{\mathbf{y}}^2 + \epsilon_2}. \tag{2}$$

(Here, ϵ_2 is the SSIM "stability coefficient".) The function $\sqrt{T(\mathbf{x}, \mathbf{y})}$ is therefore an inverse variance-weighted L^2 distance between $\mathbf{x}$ and its affine approximation $\mathbf{y}$. (It is also a metric, in fact, a *normalized metric* [4].) This is how the "unfair advantage" of flatter blocks is lessened – their generally low L^2 approximation errors are increased by division by small denominators. For more details, the reader is referred to [3].

2 Image Self-similarity in the Wavelet Domain

We consider an image of size $2^K \times 2^K$ $(K \geq 0)$ and standard tensor-product (real) wavelet basis expansions of this image [7]. The wavelet coefficients can be arranged in a standard matrix form as shown in [7]. At level k of this decomposition, $0 \leq k \leq K$, the coefficients can further be partitioned into *horizontal*, *vertical*, and *diagonal* detail blocks, denoted by $\mathbf{A}_k^h, \mathbf{A}_k^v$, and $\mathbf{A}_k^d$, respectively. Each of these blocks contains 2^{2k} wavelet coefficients a_{kij}^λ. Let A_{kij}^λ denote the unique quadtree rooted at a_{kij}^λ, $\lambda \in \{h, v, d\}$.

The self-similarity of images in the wavelet domain will be explored by examining how well "range" quadtrees in the wavelet expansion of an image are approximated by other "domain" quadtrees under affine transformations, both at the same and different scales. These approximations will have the form

$$A_{kij}^\lambda \approx \alpha A_{k'i'j'}^{\lambda'}, \qquad 0 \leq k' \leq k. \tag{3}$$

Note that the constant coefficient β coefficient is omitted: Theoretically, these wavelet expansions can have infinite length and we wish to preserve the ℓ^2-summability of the expansion coefficient trees. There are three possible cases:

- **Case 1: Purely translational.** Here, $k = k'$ and $\alpha = 1$. Quadtrees at the same level are compared with no scaling.
- **Case 3: Affine, same-scale.** Here $k = k'$ and we optimize over α. Quadtrees at the same level are compared with scaling.
- **Case 4: Affine, cross-scale.** Here, $k' < k$ and we optimize over α. Quadtrees at higher levels (hence higher resolution) are approximated by scaled copies of quadtrees from lower levels.

Each of these three cases may be viewed as a wavelet analogue of the corresponding pixel-based case discussed in the previous section. Note that Case 2 is missing since we do not consider greyscale shifts. The term "analogue" actually becomes "equivalence" in the case of the Haar wavelet basis. Proceeding in the same way as was done by Davis [6], it may be shown that for the Haar case, the pixel-based affine transformations in Cases 1, 3 and 4 induce the corresponding wavelet-based scaling transformations. (The β greyscale shifts are absorbed into higher level scaling coefficients.) We also mention that Case 4 is the basis of fractal-wavelet image coding [12].

The solution for optimal α-values for the approximation in Eq. (3) is very simple. If we let $\mathbf{a}$ and $\mathbf{c}$ denote N-vectors containing the corresponding wavelet detail coefficients from the A^{λ}_{kij} and $A^{\lambda'}_{k'i'j'}$ quadtrees, respectively, then the optimal α-value for the ℓ^2 (least-squares) approximation, $a \approx \alpha c$, is given by $\alpha = \langle \mathbf{a}, \mathbf{c} \rangle / \langle \mathbf{c}, \mathbf{c} \rangle$. Here, $\langle \cdot, \cdot \rangle$ denotes the dot product in $\mathbb{R}^N$. The associated squared ℓ^2 approximation error is given by

$$\Delta^2 \;=\; \|\mathbf{a} - \alpha \mathbf{c}\|_2^2 \;=\; \langle \mathbf{a}, \mathbf{a} \rangle - \frac{\langle \mathbf{a}, \mathbf{c} \rangle^2}{\langle \mathbf{c}, \mathbf{c} \rangle} = \|\mathbf{a}\|_2^2 \, \sin^2 \theta, \tag{4}$$

where θ is the angle between the N-vectors $\mathbf{a}$ and $\mathbf{c}$. This, of course, implies that the error associated with the approximation in Eq. (3) is bounded as follows,

$$\Delta^{\lambda,\lambda'}_{kij,k'i'j'} \;=\; \|A^{\lambda}_{kij} - \alpha A^{\lambda'}_{k'i'j'}\|_2 \;\leq\; \|A^{\lambda}_{kij}\|_2. \tag{5}$$

We see that the ℓ^2 norm of the quadtree block A^{λ}_{kij}, equivalently the square root of its energy, is the wavelet-based analogue of variance for pixels: *Wavelet coefficient quadtrees with lower energy are more approximable by the scaling transformation in Eq. (3).* We shall return to this idea in the next section.

In the experiments reported below, normalized images were once again employed, i.e., the greyscale range is $R_g = [0, 1]$. We considered $k = 6$ (there are three decompositions in total) and for Case 4, $k' = 3$. Furthermore, in order to reduce the computational cost, only quadtrees with the same orientation are compared, i.e., $\lambda = \lambda'$. The Daubechies-6 wavelet system was employed here. We mention that similar results were obtained with other systems.

Results of Numerical Experiments on Test Images

Cases 1 and 3: In Fig. 1 are shown the (same scale) Δ-error distributions of the $k = 6$ wavelet coefficient quadtrees for the normalized *Lena* and *Mandrill* images. As seen in the pixel case [2], the Case 3 distributions show greater peaking near zero error for both images. In addition, the Δ-error distributions of the *Lena* image are much more peaked near zero than for the *Mandrill* image, as was the case for pixel blocks.

We simply mention here that the histogram distributions of the α parameters associated with the Case 3 approximations for *Lena* and *Mandrill* images are remarkably similar, showing strong peaking at zero [10]. This implies that

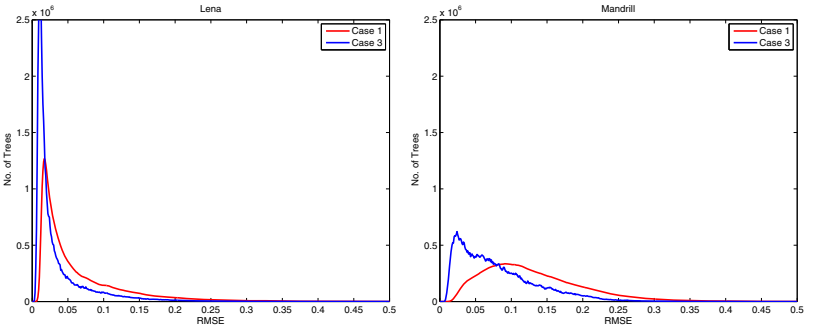

Fig. 1. Case 1 and 3 wavelet-based Δ-error distributions for *Lena* and *Mandrill*

many of the wavelet quadtrees are mapped to near zero quadtrees by the scaling transformation in Eq. (3).

Case 4: These distributions for *Lena* and *Mandrill* are also found to be virtually identical to their Case 3 counterparts in Fig. 1 [10]. We mention again that the Case 4 self-similarity exhibited by wavelet quadtrees forms the basis of fractal-wavelet coding [12]. A multiparent fractal-wavelet denoising scheme was investigated in [1].

Other test images: The Case 1,3 and 4 wavelet quadtree Δ-error distributions have been examined for all of the test images studied in [2] and [3]. For a given Case, these distributions are quite similar in shape to their pixel-based counterparts. Furthermore, the Case 4 Δ-error distribution of each image is virtually identical to its Case 3 counterpart. The degree of peaking for Case 3 distributions can be used to arrange these images in order of decreasing self-similarity, as was done in [2] from the pixel domain. The order is preserved.

Self-similarity vs Approximability in the Wavelet Domain

As discussed in Section 1, the degree of self-similarity of an image is a direct consequence of the approximability of its subblocks. In the pixel case, the approximability of a block is determined by its "flatness". Therefore, images with a higher proportion of low-variance blocks will exhibit a great degree of self-similarity, i.e., peaking of Δ-error distributions near zero.

With particular reference to Eq. (5), quadtrees with low norm/energy are more easily approximated by the scaling operation in Eq. (3). In Fig. 2 are presented histogram distributions of the norms of the $k = 6$ quadtrees for the *Lena* and *Mandrill* images. They show that, as expected, the *Lena* image contains a greater proportion of quadtrees of low norm/energy. In fact, as expected, there is a very strong similarity between these histograms and the Case 3 Δ-error distributions in Fig. 1. In the pixel domain [2], the Case 3 Δ-error distributions were similar to the distributions of standard deviations of pixel subblocks.

One can push this idea of approximability of quadtrees further, as was done in the pixel domain in [2], by using wavelet coefficient quadtrees of one image to approximate the quadtrees of another. We simply report here that results similar

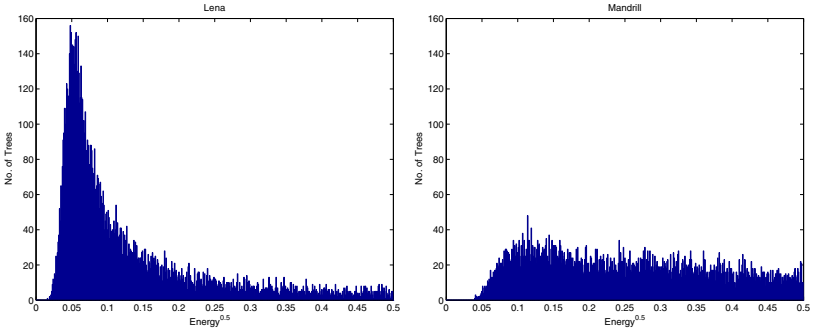

Fig. 2. Histogram distributions of norms $\|A^{\lambda}_{kij}\|_2$ of wavelet quadtrees, $k = 6$, in *Lena* and *Mandrill* images

to those found in [2] for the pixel domain have been obtained, namely, that when wavelet coefficient quadtrees of an image "Y" are used to approximate the quadtrees of image "X", the resulting Δ-error distributions are virtually identical to those obtained when the quadtrees of image "X" are used to approximate its own quadtrees [10].

Finally, it is worth mentioning that the wavelet-based results are in agreement with the pixel-based results reported earlier for the following reason: If the variance of an image subblock $u(R_i)$ is high/low, then the energy of the wavelet quadtree expansion supported on R_i is high/low.

3 The Structural Similarity (SSIM) Index

As in [3], we first express the SSIM index between two pixel blocks as a product of two components that measure (i) the similarities of their mean values and (ii) their correlation and contrast distortion. Given two non-negative signals $\mathbf{x}, \mathbf{y} \in \mathbb{R}^N_+$, the SSIM index between $\mathbf{x}$ and $\mathbf{y}$ is defined as follows,

$$S(\mathbf{x}, \mathbf{y}) = \left[\frac{2\bar{\mathbf{x}}\bar{\mathbf{y}} + \epsilon_1}{\bar{\mathbf{x}}^2 + \bar{\mathbf{y}}^2 + \epsilon_1}\right] \left[\frac{2s_{\mathbf{xy}} + \epsilon_2}{s^2_{\mathbf{x}} + s^2_{\mathbf{y}} + \epsilon_2}\right], \tag{6}$$

where

$$\bar{\mathbf{x}} = \frac{1}{N}\sum_{k=1}^{N} x_k, \qquad s_{\mathbf{xy}} = \frac{1}{N-1}\sum_{k=1}^{N}(x_k - \bar{\mathbf{x}})(y_k - \bar{\mathbf{y}}) \tag{7}$$

and $s^2_{\mathbf{x}}$ is, of course, obtained by setting $\mathbf{x} = \mathbf{y}$. The parameters ϵ_1 and ϵ_2 are small constants (relative to the maximum size of the intensities) used to provide numerical stability and accommodate the human visual system (Weber's law).

The SSIM index is symmetric: $S(\mathbf{x}, \mathbf{y}) = S(\mathbf{y}, \mathbf{x})$. Furthermore, it is bounded: $-1 \leq S(\mathbf{x}, \mathbf{y}) \leq 1$, with $S(\mathbf{x}, \mathbf{y}) = 1$ if and only if $\mathbf{x} = \mathbf{y}$. In practice, the SSIM index between two images is often computed pixelwise by using $n \times n$-pixel blocks centered at the correponding pixels of the two images. These values may be used to produce a *SSIM index map* or averaged to produce a single SSIM value.

4 A Wavelet-Based Version of the SSIM Function

We now construct a wavelet-based version of the SSIM index by expressing the components of $S(\mathbf{x}, \mathbf{y})$ in terms of wavelet coefficients [10]. In what follows, let $\mathbf{x}, \mathbf{y} \in \mathbb{R}^M$, where $M = 2^K$ for some $K \geq 0$, represent our images in the pixel domain. Now consider the projections of $\mathbf{x}$ and $\mathbf{y}$ onto a set of N orthonormal zero-mean (real) wavelet basis functions $\{\psi_k \mid 1 \leq k \leq N\}$ which correspond to a quadtree A^{λ}_{kij} with N nodes in the wavelet decompositions of $\mathbf{x}$ and $\mathbf{y}$. For example, if the quadtree is rooted at the very top of the coefficient pyramid, $N = 2^K - 1$. This set of functions spans an N-dimensional subspace of $\mathbb{R}^M$. The best L^2-based approximations of $\mathbf{x}$ and $\mathbf{y}$ in this orthonormal basis will be denoted as follows,

$$P\mathbf{x} = \sum_{k=1}^{N} a_k \psi_k, \qquad P\mathbf{y} = \sum_{k=1}^{N} c_k \psi_k, \tag{8}$$

where a_k and c_k denote the wavelet expansion coefficients of $P\mathbf{x}$ and $P\mathbf{y}$ in this basis. It now remains to express the quantities $\overline{P\mathbf{x}}$ (mean), $s_{P\mathbf{x}P\mathbf{y}}$ (covariance) and $s^2_{P\mathbf{x}}$ (variance) in terms of these expansion coefficients. These quantities will involve summations over the spatial components of $P\mathbf{x}$ and $P\mathbf{y}$. The means are trivially zero, i.e., $\overline{P\mathbf{x}} = \overline{P\mathbf{y}} = 0$, since the wavelet basis functions are zero-mean in the space/pixel coordinate. To compute the covariance, let $M' \leq M$ be the number of pixels that comprise the union of the supports of the wavelet basis functions $\psi_k, 1 \leq k \leq N$. Then enumerate these pixels so that $P\mathbf{x}$ and $P\mathbf{y}$ become M'-vectors with components $P\mathbf{x}[i]$ and $P\mathbf{y}[i]$, $1 \leq i \leq M'$, respectively. Then

$$\begin{aligned} s_{P\mathbf{x}P\mathbf{y}} &= \frac{1}{M'-1} \sum_{i=1}^{M'} (P\mathbf{x}[i] - \overline{P\mathbf{x}})(P\mathbf{y}[i] - \overline{P\mathbf{y}}) \\ &= \frac{1}{M'-1} \langle P\mathbf{x}, P\mathbf{y} \rangle \\ &= \frac{1}{M'-1} \sum_{k=1}^{N} a_k c_k. \end{aligned} \tag{9}$$

Here, $\langle \cdot, \cdot \rangle$ denotes the scalar product in $\mathbb{R}^{M'}$ and the last line follows from the orthonormality of the wavelet basis functions. Now set $\mathbf{y} = \mathbf{x}$ to obtain

$$s^2_{P\mathbf{x}} = \frac{1}{M'-1} \sum_{k=1}^{N} a_k^2. \tag{10}$$

We now substitute these results into the SSIM function in Eq. (6). The first of the two terms in Eq. (6) is 1, even in the case that $\epsilon_1 = 0$. As a result, we have the following SSIM function of the wavelet coefficient vectors $\mathbf{a}$ and $\mathbf{c}$,

$$S_W(\mathbf{a}, \mathbf{c}) = \frac{2\langle \mathbf{a}, \mathbf{c} \rangle + C_2}{\|\mathbf{a}\|_2^2 + \|\mathbf{c}\|_2^2 + C_2}, \tag{11}$$

where $C_2 = (M' - 1)\epsilon_2$.

5 Wavelet-Based Structural Self-similarity

We now examine how well quadtrees in the wavelet expansion of an image are approximated by other quadtrees according to the SSIM index, both at the same and different scales. These approximations will once again have the form in Eq. (3) – only Cases 1 and 3 (same scale) will be discussed here. However, we now optimize these approximations in terms of the SSIM index.

In order to find optimal α in the SSIM-based approximation $P\mathbf{x} = \alpha P\mathbf{y}$ or, equivalently, $\mathbf{a} = \alpha\mathbf{c}$, we must maximize the expression,

$$S_W(\mathbf{a}, \alpha\mathbf{c}), \tag{12}$$

with respect to α. For the case of zero stability parameter, $C_2 = 0$, the optimal coefficient α_{SSIM} is found to be

$$\alpha_{SSIM} = \text{sgn}(\langle \mathbf{a}, \mathbf{c} \rangle) \frac{\langle \mathbf{a}, \mathbf{a} \rangle}{\langle \mathbf{c}, \mathbf{c} \rangle}. \tag{13}$$

In the calculations reported below, wavelet quadtrees of normalized 8-bit images with the same orientation are compared, i.e., $\lambda = \lambda'$. The parameters from the L^2-based cases were used: $k = 6$ for Cases 1 and 3 and $k' = 3$ for Case 4. Once again, the Daubechies-6 wavelet system was employed.

Instead of presenting distributions of SSIM measures between quadtrees A^{λ}_{kij} and their approximations $\alpha_{SSIM} A^{\lambda'}_{k'i'j'}$, as was done in [3], we shall consider the following quantities,

$$T_W(A^{\lambda}_{kij}, \alpha_{SSIM} A^{\lambda}_{k'i'j'}) = 1 - S_W(A^{\lambda}_{kij}, \alpha_{SSIM} A^{\lambda}_{k'i'j'}). \tag{14}$$

Recall, from Eq. (1), that $T(\mathbf{x}, \mathbf{y}) = 0$ implies $\mathbf{x} = \mathbf{y}$, so that Eq. (14) represents a kind of SSIM-based approximation error. Since the range of the wavelet-based SSIM function S_W is $[-1, 1]$, it follows that the range of T_W is $[0, 2]$.

Fig. 3 shows the histogram distributions of the Case 1 and Case 3 SSIM error functions T_W of Eq. (14) for the *Lena* and *Mandrill* images. Perhaps the first noteworthy feature is that the Case 3 (blue) distributions are supported only on the interval $[0, 1]$. This is because the SSIM value of the optimal Case 3 SSIM-based approximation $\mathbf{x} \approx \alpha\mathbf{y}$ is non-negative [3]. This property contributes to the fact that the Case 3 (blue) distributions have a greater concentration at low-to-zero error than their Case 1 (red) counterparts. Once again, an improvement is expected since Case 3 involves an optimization over an α parameter, whereas there is no optimization in Case 1. That being said, there is much less improvement for the *Mandrill* image than the *Lena* image.

The peaking of the Case 3 histograms is certainly not as dramatic as in the L^2-based histograms of Fig. 1. In other words, the *Lena* and *Mandrill* do not exhibit a high degree of self-similarity in the wavelet domain in terms of SSIM. That being said, one can still claim that the *Lena* image is more self-similar than the *Mandrill* image in terms of SSIM.

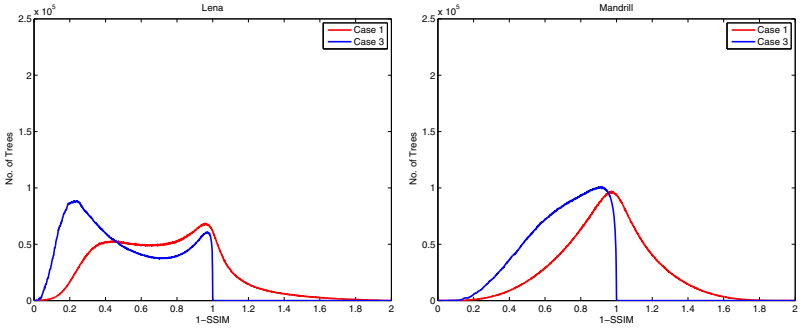

Fig. 3. Case 1 (red) and Case 3 (blue) distributions of the wavelet-based SSIM error $T_W = 1 - S_W$ *Lena* and *Mandrill* images

Finally, we mention that for the other test images employed in the L^2-based study [2], the results were consistent: In the wavelet domain, images are not as self-similar with respect to SSIM as they are with respect to ℓ^2. This can once again be explained in terms of normalized metrics. From Eqs. (14) and (11),

$$T_W(\mathbf{a}, \mathbf{c}) = \frac{\|\mathbf{a} - \mathbf{c}\|_2^2}{\|\mathbf{a}\|_2^2 + \|\mathbf{c}\|_2^2 + C_2}. \tag{15}$$

This shows that $T_W(\mathbf{a}, \mathbf{c})$ is an inverse energy-weighted squared ℓ^2 distance – the wavelet analogue of Eq. (2). The "unfair advantage" of quadtrees of low energy is lessened in the same way as in the pixel domain: their ℓ^2 distances are increased by division by small denominators.

6 Concluding Remarks

We have demonstrated that the self-similarity of an image in the wavelet domain with respect to the ℓ^2 distance is determined by the proportion of wavelet coefficient quadtrees with low energy/norm. The high degree of self-similarity demonstrated by images is reduced when the SSIM index is used to compare quadtrees. The SSIM-based error measure T_W, an inverse energy-weighted ℓ^2 distance, reduces the "unfair advantage" of quadtrees with low energy/norm. This is a perfect analogue to self-similarity in the pixel domain.

Perhaps some final comments are in order regarding the potential use of self-similarity in the wavelet domain. One practical application of Case 4 self-similarity in the wavelet domain is the fractal-wavelet transform [12] which is useful in image compression [6], image super-resolution and image denoising [9]. On the other hand, we have investigated several wavelet-based nonlocal means denoising schemes but with little success. The tails of quadtrees of noisy images are dominated by the noise and a nonlocal means strategy will not be effective in such situations with virtually zero signal-to-noise ratio.

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 (DG).

References

1. Alexander, S.K.: Multiscale Methods in Image Modelling and Image Processing, Ph.D. Thesis, Department of Applied Mathematics, University of Waterloo (2005)
2. Alexander, S.K., Vrscay, E.R., Tsurumi, S.: A Simple, General Model for the Affine Self-similarity of Images. In: Campilho, A., Kamel, M.S. (eds.) ICIAR 2008. LNCS, vol. 5112, pp. 192–203. Springer, Heidelberg (2008)
3. Brunet, D., Vrscay, E.R., Wang, Z.: Structural Similarity-Based Affine Approximation and Self-similarity of Images Revisited. In: Kamel, M., Campilho, A. (eds.) ICIAR 2011, Part II. LNCS, vol. 6754, pp. 264–275. Springer, Heidelberg (2011)
4. Brunet, D., Vrscay, E.R., Wang, Z.: On the mathematical properties of the structural similarity index. To appear in IEEE Transactions on Image Processing
5. 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); an updated and expanded version of this paper appears in SIAM Review 52, 113–147 (2010)
6. Davis, G.: A wavelet-based analysis fractal image compression. IEEE Trans. Image Proc. 7, 141–154 (1998)
7. Daubechies, I.: Ten Lectures on Wavelets. SIAM Press, Philadelphia (1992)
8. Fisher, Y. (ed.): Fractal Image Compression: Theory and Application. Springer, New York (1995)
9. Ghazel, M., Freeman, G., Vrscay, E.R.: Fractal-wavelet image denoising revisited. IEEE Trans. Image Proc. 15(9), 2669–2675 (2006)
10. Glew, D.: Self-similarity of images and nonlocal image processing, M. Math. Thesis, University of Waterloo (2011)
11. Lu, N.: Fractal Imaging. Academic Press, New York (1997)
12. Vrscay, E.R.: A generalized class of fractal-wavelet transforms for image representation and compression. Can. J. Elec. Comp. Eng. 23, 69–83 (1998)
13. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: From error visibility to structural similarity. IEEE Trans. Image Proc. 13, 600–612 (2004)
14. Wang, Z., Bovik, A.C.: Mean squared error: Love it or leave it? A new look at signal fidelity measures. IEEE Sig. Proc. Mag. 26, 98–117 (2009)