

Structural similarity-based approximation of signals and images using orthogonal bases

Dominique Brunet¹, Edward R. Vrscay¹, and Zhou Wang²

¹ Department of Applied Mathematics, Faculty of Mathematics, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1

² Department of Electrical and Computer Engineering, Faculty of Engineering, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1
dbrunet@uwaterloo.ca, ervrscay@uwaterloo.ca, zhouwang@ieee.org

Abstract. The structural similarity (SSIM) index has been shown to be an useful tool in a wide variety of applications that involve the assessment of image quality and similarity. However, in-depth studies are still lacking on how to incorporate it for signal representation and approximation problems, where minimal mean squared error is still the dominant optimization criterion. Here we examine the problem of best approximation of signals and images by maximizing the SSIM between them. In the case of a decomposition of a signal in terms of an orthonormal basis, the optimal SSIM-based coefficients are determined with a surprisingly simple approach, namely, a scaling of the optimal L^2 coefficients. We then examine a very simple algorithm to maximize SSIM with a constrained number of basis functions. The algorithm is applied to the DCT approximation of images.

1 Introduction

The structural similarity (SSIM) index [9] was proposed as a measure to predict visual distortions between two images. If one of the images being compared is assumed to have perfect quality, the SSIM value can also be interpreted as a perceptual quality measure of the second image. When tested with large-scale independent subject-rated image quality databases [6, 4], SSIM has demonstrated superior performance in comparison with traditional image distortion measures such as the mean square error (MSE), which is the most widely employed metric in the image processing literature [8]. In the past few years, SSIM has found a wide range of applications, ranging from image compression, restoration, fusion, and watermarking, to video streaming, digital camera design, biometrics, remote sensing and target recognition [8]. In most of the existing works, however, SSIM has been used for quality evaluation and algorithm comparison purposes only. Much less has been done on using SSIM as an optimization criterion in the design and tuning of image processing algorithms and systems [10, 3, 1, 2, 5].

A fundamental issue that has to be resolved before effectively employing SSIM-based optimization in various image processing applications is how to decompose a signal or image as a linear combination of basis functions optimally *in*

the SSIM sense (as opposed to the usual L^2 sense). This is a nontrivial problem, given the non-convex property of SSIM. It has been addressed in the particular contexts of image compression [3, 5] and image restoration [1, 2]. In this paper, however, we “step back” in an effort to understand the problem mathematically. We analyze the simpler case where the SSIM function is defined over nonoverlapping blocks (as opposed to “local SSIM,” which involves overlapping patches) and an orthonormal basis is employed. The unique global maximum of the SSIM function over a block may be found by examining its partial derivatives with respect to the expansion coefficients. In this way, an optimal SSIM-based approximation is obtained, as opposed to the well-known L^2 -based result, i.e., mean-squared-error (MSE). We obtain the remarkable result that the optimal SSIM-based approximation may, in fact, be determined from the optimal L^2 -based approximation: The zeroth-order coefficients are the same, and the higher order SSIM coefficients are obtained from their Fourier counterparts by scaling.

In closing this section, we mention that partial derivatives of the SSIM have been used before. In fact, formulas for the more complicated “local SSIM” case appear in [11]. In that paper, they were used in a numerical gradient ascent/descent algorithm for finding the maxima and minima of the SSIM function over spheres of constant MSE with respect to a reference image. In this case, however, the formulas for the derivatives are very complicated and stationary points cannot, in general, be determined analytically. The study presented below is intended to be a first step toward a deeper understanding of the relationship between SSIM- and L^2 -based approximations.

2 SSIM-based approximations of signals/images

Very briefly, the “local SSIM,” that is, the SSIM computed between two local image patches, say \mathbf{a} and \mathbf{b} , measures the similarities of three elements of these patches: (i) the local patch luminances or brightness values, (ii) the local patch contrasts and (iii) local patch structures. These three components are then multiplied to form a local SSIM index between \mathbf{a} and \mathbf{b} . The “closer” that \mathbf{a} and \mathbf{b} are to each other, the closer the value of the SSIM to 1. It is possible to express this SSIM function as a product of only two components. It is this form of the SSIM that is employed in this paper.

In what follows, we let $\mathbf{x}, \mathbf{y} \in \mathbf{R}^N$ denote two N -dimensional signal/image blocks, e.g., $\mathbf{x} = (x_1, x_2, \dots, x_N)$. We consider a variation of the SSIM function described above, a global SSIM between \mathbf{x} and \mathbf{y} defined as follows,

$$S(\mathbf{x}, \mathbf{y}) = S_1(\mathbf{x}, \mathbf{y})S_2(\mathbf{x}, \mathbf{y}) = \left[\frac{2\bar{\mathbf{x}}\bar{\mathbf{y}} + \epsilon_1}{\bar{\mathbf{x}}^2 + \bar{\mathbf{y}}^2 + \epsilon_1} \right] \left[\frac{2s_{\mathbf{x}\mathbf{y}} + \epsilon_2}{s_{\mathbf{x}}^2 + s_{\mathbf{y}}^2 + \epsilon_2} \right], \quad (1)$$

where

$$\bar{\mathbf{x}} = \frac{1}{N} \sum_{i=1}^N x_i, \quad s_{\mathbf{x}}^2 = \frac{1}{N-1} \sum_{i=1}^N (x_i - \bar{\mathbf{x}})^2, \quad s_{\mathbf{x}\mathbf{y}} = \frac{1}{N-1} \sum_{i=1}^N (x_i - \bar{\mathbf{x}})(y_i - \bar{\mathbf{y}}). \quad (2)$$

The small positive constants $\epsilon_1, \epsilon_2 \ll 1$ are added for numerical stability and can be adjusted to accommodate the perception of the human visual system. It will be convenient to denote the special case where these parameters are zero as follows,

$$S_0(\mathbf{x}, \mathbf{y}) = \frac{4\bar{\mathbf{x}}\bar{\mathbf{y}}s_{\mathbf{x}\mathbf{y}}}{(\bar{\mathbf{x}}^2 + \bar{\mathbf{y}}^2)(s_{\mathbf{x}}^2 + s_{\mathbf{y}}^2)}. \quad (3)$$

The functional form of the component S_1 in Eq. (1) was originally chosen in an effort to accommodate Weber's law of perception [9]; that of S_2 follows the idea of divisive normalization [7].

Note that $-1 \leq S(\mathbf{x}, \mathbf{y}) \leq 1$, and $S(\mathbf{x}, \mathbf{y}) = 1$ if and only if $\mathbf{x} = \mathbf{y}$. The component $S_1(\mathbf{x}, \mathbf{y})$ measures the similarity between the means of \mathbf{x} and \mathbf{y} : If $\bar{\mathbf{x}} = \bar{\mathbf{y}}$, then $S_1(\mathbf{x}, \mathbf{y}) = 1$, its maximum possible value. This will be important in the discussion below.

Unless otherwise specified, we consider \mathbf{x} to be a given signal and \mathbf{y} to be an approximation to \mathbf{x} . We shall generally consider \mathbf{y} to be an element of a particular subset $A \subset \mathbf{R}^N$ – details to be given below – and look for solutions to the problem

$$\mathbf{y}_A = \arg \max_{\mathbf{y} \in A} S(\mathbf{x}, \mathbf{y}). \quad (4)$$

In the case $A = \mathbf{R}^N$, $\mathbf{y} = \mathbf{x}$ and $S(\mathbf{x}, \mathbf{y}) = S(\mathbf{x}, \mathbf{x}) = 1$.

We start with a set of (complete) orthonormal basis functions \mathbf{R}^N , to be denoted as $\{\psi_0, \psi_1, \dots, \psi_{N-1}\}$. We assume that only the first element has nonzero mean: $\bar{\psi}_k = 0$ for $1 \leq k \leq N-1$. We also assume that ψ_0 is “flat”, i.e., constant: $\psi_0 = N^{-1/2}(1, 1, \dots, 1)$, which accommodates the discrete cosine transform (DCT) as well as Haar multiresolution system on \mathbf{R}^N .

The L^2 -based expansion of \mathbf{x} in this basis is, of course,

$$\mathbf{x} = \sum_{k=0}^{N-1} a_k \psi_k, \quad a_k = \langle \mathbf{x}, \psi_k \rangle, \quad 0 \leq k \leq N-1. \quad (5)$$

It follows that

$$\bar{\mathbf{x}} = a_0 N^{-1/2}. \quad (6)$$

The expansions of the approximation \mathbf{y} will be denoted as follows,

$$\mathbf{y} = \mathbf{y}(\mathbf{c}) = \sum_{k=0}^{N-1} c_k \psi_k, \quad (7)$$

where the notation $\mathbf{y}(\mathbf{c})$ acknowledges the dependence of the approximation on the coefficients c_k . It also follows that

$$\bar{\mathbf{y}} = c_0 N^{-1/2}. \quad (8)$$

In this study, the approximation spaces A in (4) will be the spans of subsets of the set of basis functions $\{\psi_k\}_{k=0}^{N-1}$ which include ψ_0 : From (6) and (8), the

inclusion of ψ_0 automatically maximizes component $S_1(\mathbf{x}, \mathbf{y})$ of the SSIM function. At this point, we do not specify exactly which other ψ_k basis functions will be used, but consider all possible subsets of $M < N$ basis functions:

$$A = \text{span}\{\psi_0, \psi_{\gamma(1)}, \dots, \psi_{\gamma(M-1)}\}, \quad (9)$$

where $\gamma(k) \in \{1, 2, \dots, N-1\}$ and $c_{\gamma(M)} = \dots = c_{\gamma(N-1)} = 0$. Of course, we are interested in finding the *optimal* M -dimensional subset, in the SSIM sense.

Before studying optimal SSIM approximations, however, it is most helpful to review the well-known L^2 -based case.

Proposition 1. *For a given $\mathbf{x} \in \mathbf{R}^N$, the M coefficients c_k of the optimal L^2 -based approximation $\mathbf{y} \in A$ to \mathbf{x} are given by $c_0 = \langle \mathbf{x}, \psi_0 \rangle$ and the $M-1$ remaining Fourier coefficients $a_k = \langle \mathbf{x}, \psi_k \rangle$ of greatest magnitude, i.e.,*

$$c_k = a_{\gamma(k)} = \langle \mathbf{x}, \psi_{\gamma(k)} \rangle, \quad 1 \leq k \leq M-1, \quad (10)$$

where $|a_{\gamma(1)}| \geq |a_{\gamma(2)}| \geq \dots \geq |a_{\gamma(M-1)}| \geq |a_l|$ with $l \in \{1, 2, \dots, N-1\} \setminus \{\gamma(1), \dots, \gamma(M-1)\}$.

Proof. For an arbitrary $\mathbf{c} \in \mathbf{R}^N$, let $\Delta(\mathbf{x}, \mathbf{y}(\mathbf{c})) = \|\mathbf{x} - \mathbf{y}(\mathbf{c})\|_2$, the L^2 error of approximation of \mathbf{x} by $\mathbf{y}(\mathbf{c})$. For any $p \in \{0, 1, 2, \dots, N-1\}$, consider the change in this error produced by altering the coefficient c_p by ϵ , i.e., $\mathbf{c} \rightarrow \mathbf{c} + \epsilon \hat{\mathbf{e}}_p$. Because the squared L^2 -error is a quadratic form in the c_k , its Taylor series in ϵ ,

$$\Delta^2(\mathbf{x}, \mathbf{y}(\mathbf{c} + \epsilon \hat{\mathbf{e}}_p)) = \Delta^2(\mathbf{x}, \mathbf{y}(\mathbf{c})) + \epsilon \frac{\partial}{\partial c_p} \Delta^2(\mathbf{x}, \mathbf{y}(\mathbf{c})) + \dots, \quad (11)$$

is finite – in fact, a quadratic polynomial:

$$\Delta^2(\mathbf{x}, \mathbf{y}(\mathbf{c} + \epsilon \hat{\mathbf{e}}_p)) = \|\mathbf{x} - \mathbf{y}(\mathbf{c})\|_2^2 + 2\epsilon(c_p - \langle \mathbf{x}, \psi_p \rangle) + \epsilon^2. \quad (12)$$

We see that the only stationary point occurs when $c_p = a_p = \langle \mathbf{x}, \psi_p \rangle$ and that it is, in fact, a global minimum. Since ψ_0 is a basis element of A , we set $c_0 = a_0 = \langle \mathbf{x}, \psi_0 \rangle$. Now, for some choice of distinct $M-1$ indices $\gamma(k) \in \{1, 2, \dots, N-1\}$, set $c_k = a_{\gamma(k)}$, $1 \leq k \leq M-1$, and $c_l = 0$ otherwise. Then

$$\begin{aligned} \|\mathbf{x} - \mathbf{y}(\mathbf{c})\|_2^2 &= \|\mathbf{x}\|_2^2 - 2 \sum_{k=0}^{N-1} c_k \langle \mathbf{x}, \psi_k \rangle + \sum_{k=0}^{N-1} c_k^2 \\ &= \|\mathbf{x}\|_2^2 - a_0^2 - \sum_{k=1}^{M-1} a_{\gamma(k)}^2. \end{aligned} \quad (13)$$

Clearly, the smallest L^2 approximation error is produced if the $M-1$ coefficients $a_{\gamma(k)}$ in Eq. (13) are those with the largest magnitudes. \square

We now consider the optimal SSIM-based approximation of an element $\mathbf{x} \in \mathbf{R}^N$ in the M -dimensional subspace A defined in Eq. (9).

Proposition 2. *The coefficients of the optimal SSIM-based approximation of \mathbf{x} in the M -dimensional subspace A defined in Eq. (9) are given by*

$$c_0 = a_0, \quad c_{\gamma(k)} = \alpha a_{\gamma(k)} \quad \text{for } 1 \leq k \leq M-1, \quad (14)$$

where the $a_{\gamma(k)} = \langle x, \psi_{\gamma(k)} \rangle$ are the $M-1$ optimal L^2 -based coefficients from Proposition 1 and the scaling coefficient α is given by

$$\alpha = \frac{-\epsilon_2 + \sqrt{\epsilon_2^2 + \left(\frac{4}{N-1} \sum_{k=1}^{M-1} a_{\gamma(k)}^2\right)(s_{\mathbf{x}}^2 + \epsilon_2)}}{\frac{2}{N-1} \sum_{k=1}^{M-1} a_{\gamma(k)}^2}. \quad (15)$$

Proof. Without loss of generality, we assume that the last $N-M$ coefficients c_k are zeros. After some simple algebra, we find that

$$s_{\mathbf{xy}} = \frac{1}{N-1} \sum_{k=1}^{N-1} a_k c_k \quad \text{and} \quad s_{\mathbf{y}}^2 = \frac{1}{N-1} \sum_{k=1}^{N-1} c_k^2. \quad (16)$$

The dependence of the SSIM function in Eq. (1) on the c_k is as follows: (i) the first term in Eq. (1) depends only on the coefficient c_0 and (ii) the second term in Eq. (1) is independent of c_0 . The choice $c_0 = a_0$ maximizes the first term in Eq. (1), giving it the value of 1. In terms of the remaining c_k ,

$$S(\mathbf{x}, \mathbf{y}(\mathbf{c})) = \frac{\frac{2}{N-1} \sum_{k=1}^{M-1} a_k c_k + \epsilon_2}{s_{\mathbf{x}}^2 + \frac{1}{N-1} \sum_{k=1}^{M-1} c_k^2 + \epsilon_2}. \quad (17)$$

We now look for stationary points which will be candidates for solutions to the approximation problem in (4). Logarithmic differentiation yields the following partial derivatives with respect to c_k for $1 \leq k \leq M-1$:

$$\begin{aligned} \frac{\partial S}{\partial c_k} = S \left[\frac{2\bar{x}}{2\bar{x}\bar{y} + \epsilon_1} \frac{\partial \bar{y}}{\partial c_k} + \frac{2}{2s_{\mathbf{xy}} + \epsilon_2} \frac{\partial s_{\mathbf{xy}}}{\partial c_k} \right. \\ \left. - \frac{2\bar{y}}{\bar{x}^2 + \bar{y}^2 + \epsilon_1} \frac{\partial \bar{y}}{\partial c_k} - \frac{1}{s_{\mathbf{x}}^2 + s_{\mathbf{y}}^2 + \epsilon_2} \frac{\partial s_{\mathbf{y}}^2}{\partial c_k} \right]. \quad (18) \end{aligned}$$

After some additional (yet simple) algebra, we obtain the following conditions for a stationary point,

$$\frac{\partial S}{\partial c_k} = \frac{S}{N-1} \left[\frac{2a_k}{2s_{\mathbf{xy}} + \epsilon_2} - \frac{2c_k}{s_{\mathbf{x}}^2 + s_{\mathbf{y}}^2 + \epsilon_2} \right] = 0, \quad 1 \leq k \leq M-1. \quad (19)$$

If $a_p = 0$ for any $1 \leq p \leq M-1$, then $c_p = 0$. Otherwise, we have that

$$\frac{a_k}{c_k} = \frac{2s_{\mathbf{xy}} + \epsilon_2}{s_{\mathbf{x}}^2 + s_{\mathbf{y}}^2 + \epsilon_2}, \quad 1 \leq k \leq M-1. \quad (20)$$

Note that the RHS of each equation is independent of k , implying that

$$\frac{a_1}{c_1} = \frac{a_2}{c_2} = \dots = \frac{a_{M-1}}{c_{M-1}} = C \quad (\text{constant}). \quad (21)$$

Hence

$$c_k = \alpha a_k \quad \text{for } 1 \leq k \leq M-1, \quad (22)$$

where $\alpha = 1/C$. We now rewrite Eq. (20) as follows,

$$(2s_{\mathbf{x}\mathbf{y}}(\mathbf{c}) + \epsilon_2)c_k = (s_{\mathbf{x}}^2 + s_{\mathbf{y}}^2(\mathbf{c}) + \epsilon_2)a_k, \quad (23)$$

and employ (22) and the expansions in Eq. (16) to arrive at the following quadratic equation in α :

$$\frac{2\alpha^2}{N-1} \sum_{k=1}^{M-1} a_k^2 + \alpha\epsilon_2 = s_{\mathbf{x}}^2 + \frac{\alpha^2}{N-1} \sum_{k=1}^{M-1} a_k^2 + \epsilon_2. \quad (24)$$

The roots of this equation are

$$\alpha_{1,2} = \frac{-\epsilon_2 \pm \sqrt{\epsilon_2^2 + \left(\frac{4}{N-1} \sum_{k=1}^{M-1} a_k^2\right)(s_{\mathbf{x}}^2 + \epsilon_2)}}{\frac{2}{N-1} \sum_{k=1}^{M-1} a_k^2}. \quad (25)$$

Notice that $\alpha_1 \geq 1$ and $\alpha_2 \leq -1$. Substituting $c_k = \alpha_{1,2}a_k$ into Eq. (17) we observe that α_1 and α_2 correspond to the scaling factors for, respectively, a local maximum and a local minimum.

Now that the natures of the critical points have determined, we examine the behaviour of $S(\mathbf{x}, \mathbf{y})$ “on the boundaries,” i.e., as $|c_k| \rightarrow \infty$. In this case, $|S(\mathbf{x}, \mathbf{y})| \rightarrow 0$, which allows us to conclude that $c_k = \alpha_1 a_k$ at the global maximum, thus proving Eq. (15). For the remainder of the proof, we let $\alpha = \alpha_1$.

For a given $M < N$, it now remains to determine which subset of $M-1$ coefficients c_k should be chosen in order to maximize the structural similarity $S(\mathbf{x}, \mathbf{y}(\mathbf{c}))$. The global maximum value of the structural similarity, which we denote as S_{\max} , is found by substituting $c_k = \alpha a_k$, $1 \leq k \leq M-1$, into Eq. (17):

$$S_{\max}(\mathbf{x}, \mathbf{y}(\mathbf{c})) = \frac{\frac{2\alpha}{N-1} \sum_{k=1}^{M-1} a_k^2 + \epsilon_2}{s_{\mathbf{x}}^2 + \frac{\alpha^2}{N-1} \sum_{k=1}^{M-1} a_k^2 + \epsilon_2}. \quad (26)$$

From Eq. (24), we have the interesting result that

$$S_{\max}(\mathbf{x}, \mathbf{y}(\mathbf{c})) = \frac{1}{\alpha}. \quad (27)$$

Substitution of this result into Eq. (26) yields a quadratic equation in S_{\max} . Only the positive root of this equation is admissible and it is given by

$$S_{\max}(\mathbf{x}, \mathbf{y}(\mathbf{c})) = \frac{\epsilon_2 + \sqrt{\epsilon_2^2 + \left(\frac{4}{N-1} \sum_{k=1}^{M-1} a_k^2\right)(s_{\mathbf{x}}^2 + \epsilon_2)}}{2(s_{\mathbf{x}}^2 + \epsilon_2)}. \quad (28)$$

From this expression, it is clear that the maximum possible value of S_{\max} is achieved if the $M-1$ Fourier coefficients a_k with the largest magnitudes are employed in the summation. \square

A few remarks regarding these results are in order.

1. An important consequence of Proposition 2 is that the optimal SSIM-based approximation $\mathbf{y}(\mathbf{c})$ of \mathbf{x} may be obtained by first computing the best L^2 approximation of \mathbf{x} that includes ψ_0 (which is almost always the case), then setting $c_0 = a_0$ and finally scaling the remaining Fourier coefficients according to Eqs. (14) and (15).
2. Since $c_0 = a_0$, it follows that $\bar{\mathbf{x}} = \bar{\mathbf{y}}(\mathbf{c})$. Regarding the other coefficients, the fact that the scaling factor $\alpha > 1$ for $M < N$ implies that the SSIM-based approximation \mathbf{y} represents a contrast-enhanced version of \mathbf{x} .
3. In the special case that $\epsilon_2 = 0$, the optimal scaling factor α in Eq. (15) has the simple form

$$\alpha = \left[\sum_{k=1}^{N-1} a_k^2 \right]^{1/2} \left[\sum_{k=1}^{M-1} a_k^2 \right]^{-1/2}. \quad (29)$$

4. In the special case $M = N$, we have $\alpha = 1$ and $S_{\max} = 1$, as expected since there is no approximation, i.e., $\mathbf{y} = \mathbf{x}$.
5. The existence of such a simple analytic solution to this problem is made possible by the simplicity of the approach – we have been considering “global” SSIM, i.e., the entire signal/image (or block), as opposed to “local SSIM” where overlapping patches/neighbourhoods are employed. In the latter case, the derivatives of the SSIM function with respect to the c_k coefficients are not as straightforward. The above approach applies directly to (nonoverlapping) block-based coding, which includes DCT and Haar wavelet coding.
6. The assumption that the first function ψ_0 is “flat” may be relaxed, in which case Eqs. (6) and (8) would have to be modified. This, however, will not change the condition that $c_0 = a_0$ in the SSIM-optimality condition.

We finally remark that SSIM-based approximation (with no stability constants) may be viewed as a kind of (inverse) variance-weighted L^2 approximation of signals after their means have been subtracted out, as shown by Richter and Kim [5]. To see this result, let $\mathbf{x}, \mathbf{y} \in \mathbf{R}^N$ and define $\mathbf{x}_0 = \mathbf{x} - \bar{\mathbf{x}}$ and $\mathbf{y}_0 = \mathbf{y} - \bar{\mathbf{y}}$ so that $\bar{\mathbf{x}}_0 = \bar{\mathbf{y}}_0 = 0$. Then

$$\begin{aligned} \|\mathbf{x}_0 - \mathbf{y}_0\|^2 &= \sum_{k=1}^N (x_{0,k} - y_{0,k})^2 \\ &= (N-1)[s_{\mathbf{x}_0}^2 + s_{\mathbf{y}_0}^2 - 2s_{\mathbf{x}_0\mathbf{y}_0}]. \end{aligned} \quad (30)$$

From this, the definition of the SSIM function $S_0(\mathbf{x}, \mathbf{y})$ in Eq. (1) and a little algebra, we find that

$$1 - S_0(\mathbf{x}_0, \mathbf{y}_0) = \frac{1}{N-1} \frac{\|\mathbf{x}_0 - \mathbf{y}_0\|^2}{s_{\mathbf{x}_0}^2 + s_{\mathbf{y}_0}^2}. \quad (31)$$

This, along with Proposition 2, gives an idea of the link between L^2 - and SSIM-based approximations.

3 SSIM-based image reconstruction from a constrained number of DCT coefficients

Proposition 2 provides a very simple procedure to optimize L^2 -based expansions from the SSIM point of view. The procedure, however, has very limited applicability. When the the same number M of coefficients are employed in the L^2 - and SSIM-based optimizations, we find that the latter generally yields very little, if any, noticeable improvement in perceptual quality. Indeed, the greatest increase in the SSIM value is usually found for small values of M , in which case both the L^2 - and SSIM-based optimizations yield poor approximations.

On the other hand, SSIM-based optimization may yield significant improvements in perceptual quality when it is employed to decide the allocation of a prescribed number of coefficients/bits. We illustrate below with an application to the block-based discrete cosine transform (DCT).

Our simple algorithm starts with the set of zeroth-order coefficients for all blocks. The goal is to add K of the remaining higher-order DCT coefficients to this set. At each step of the selection process, we estimate the gain in structural similarity (with respect to the original image, using the original DCT coefficients) produced by adding a DCT coefficient that has not yet been employed. The unused coefficients from all blocks of the image are examined. The DCT coefficient yielding the greatest increase in the SSIM is then added to the set. For comparison, we perform a similar algorithm in which the decrease in the L^2 error of approximation is used as the criterion for selection at each step.

We define the following quantities:

1. BSSIM: the average value of the structural similarities, (Eq. 1), of all non-overlapping blocks of the image,
2. MSSIM: the average of the *weighted* SSIM, computed with a Gaussian sliding window with the parameters as in [9]. Our method will give us a BSSIM-optimal reconstructed image, but not necessarily a MSSIM-optimal image,
3. \mathbf{x}^i : the i -th block of the image \mathbf{x} being approximated,
4. $V_{K_i}^i := \frac{1}{N-1} \sum_{k=1}^{K_i} (a_{\gamma(k)}^i)^2$: the variance of the DCT approximation of the \mathbf{x}^i using K_i non-zero higher-order DCT coefficients.
5. $\mathbf{c}_{K_i}^i$: the set of K_i non-zero higher-order coefficients which, along with the zeroth-order coefficient, define the SSIM-based approximation \mathbf{y}^i to block \mathbf{x}^i .

For a given block \mathbf{x}^i , the gain in structural similarity produced by adding the first non-zero coefficient, c_1^i , is, from Eq. (28),

$$S(\mathbf{x}^i, \mathbf{y}^i(\mathbf{c}_1^i)) - S(\mathbf{x}^i, \mathbf{y}^i(\mathbf{c}_0^i)) = \frac{\sqrt{\epsilon_2^2 + 4V_1^i(s_{\mathbf{x}^i}^2 + \epsilon_2)} - \epsilon_2}{2(s_{\mathbf{x}^i}^2 + \epsilon_2)}. \quad (32)$$

For a given block \mathbf{x}^i , the gain in SSIM produced by adding the coefficient $c_{K_i+1}^i$ ($K_i > 0$) to the existing set of SSIM-optimized DCT coefficients is given by

$$S(\mathbf{x}^i, \mathbf{y}^i(\mathbf{c}_{K_i+1}^i)) - S(\mathbf{x}^i, \mathbf{y}^i(\mathbf{c}_{K_i}^i)) = \frac{\sqrt{\epsilon_2^2 + 4V_{K_i+1}^i(s_{\mathbf{x}^i}^2 + \epsilon_2)} - \sqrt{\epsilon_2^2 + 4V_{K_i}^i(s_{\mathbf{x}^i}^2 + \epsilon_2)}}{2(s_{\mathbf{x}^i}^2 + \epsilon_2)}. \quad (33)$$

After examining all blocks, the coefficient $c_{K_{i^*}+1}^{i^*}$ yielding the highest gain in SSIM, according to either Eq. (32) ($K_{i^*} = 0$) or Eq. (33) ($K_{i^*} > 0$) is then added to the set. The algorithm is terminated when K coefficients have been added.

As an example, the 512×512 8bpp *Lena* test image was decomposed with the discrete cosine transform (DCT) over 8×8 nonoverlapping pixel blocks. In Fig. 1 are shown the results for a “total budget” of $K = 10,000$ non-zero higher-order coefficients. (4096 zeroth-order c_0^i coefficients are also employed in the expansion, but not counted, in this simple scheme.) In the first row of this figure, we show (a) the original *Lena* image, (b) the “BSSIM map” of the optimal BSSIM DCT approximation and (c) the “BSSIM map” of the optimal L^2 DCT approximation that employs the same budget. In (b) and (c), the greyscale assigned to each block is proportional to its SSIM value, with black representing 0, and white representing 1.

Note that the BSSIM map in (b) is, for the most part, “lighter” than that in (c) which, of course, is expected, since the former corresponds to SSIM optimization. However, there are some blocks, most notably those containing edges, in which the SSIM values of (b) are lower than those of (c). This is revealed in the lower portion of Fig. 1, where three representative patches – (i) a part of the hat, (ii) the face and (iii) the shoulder – of the *Lena* image are presented. In the first column are shown the three patches from the original image. The second column of the array presents the patches from the SSIM-optimized image. The third column shows the patches from the L^2 -optimized image. We see that the BSSIM-optimized approximation demonstrates less blocking effects on smooth regions than its L^2 -optimized counterpart. It also preserves more details such as the fine textures of the hat – in fact, the improvement afforded by the SSIM method over the L^2 method is quite remarkable here.

On other hand, the BSSIM-optimized approximation performs less efficiently on edges. Here, the contrasts – as determined by the scaling factor α for each block – are too large, giving rise to noticeable blocking effects. This could be explained in part by the fact that the BSSIM-optimal image enhances the contrast of the image locally on each block without taking neighbouring blocks into account. An MSSIM-optimized approximation could correct such blocking problems around edges.

The excessively large contrasts exhibited in these blocks may also be due to a small number of coefficients being assigned to them in the BSSIM procedure, with its “total bit budget” for the entire image. From Proposition 2, the scaling function α in Eq. (15) is greater than one, and approaches one as the number of coefficients M approaches N , the size of the block. In the case $M = 2$ (corre-

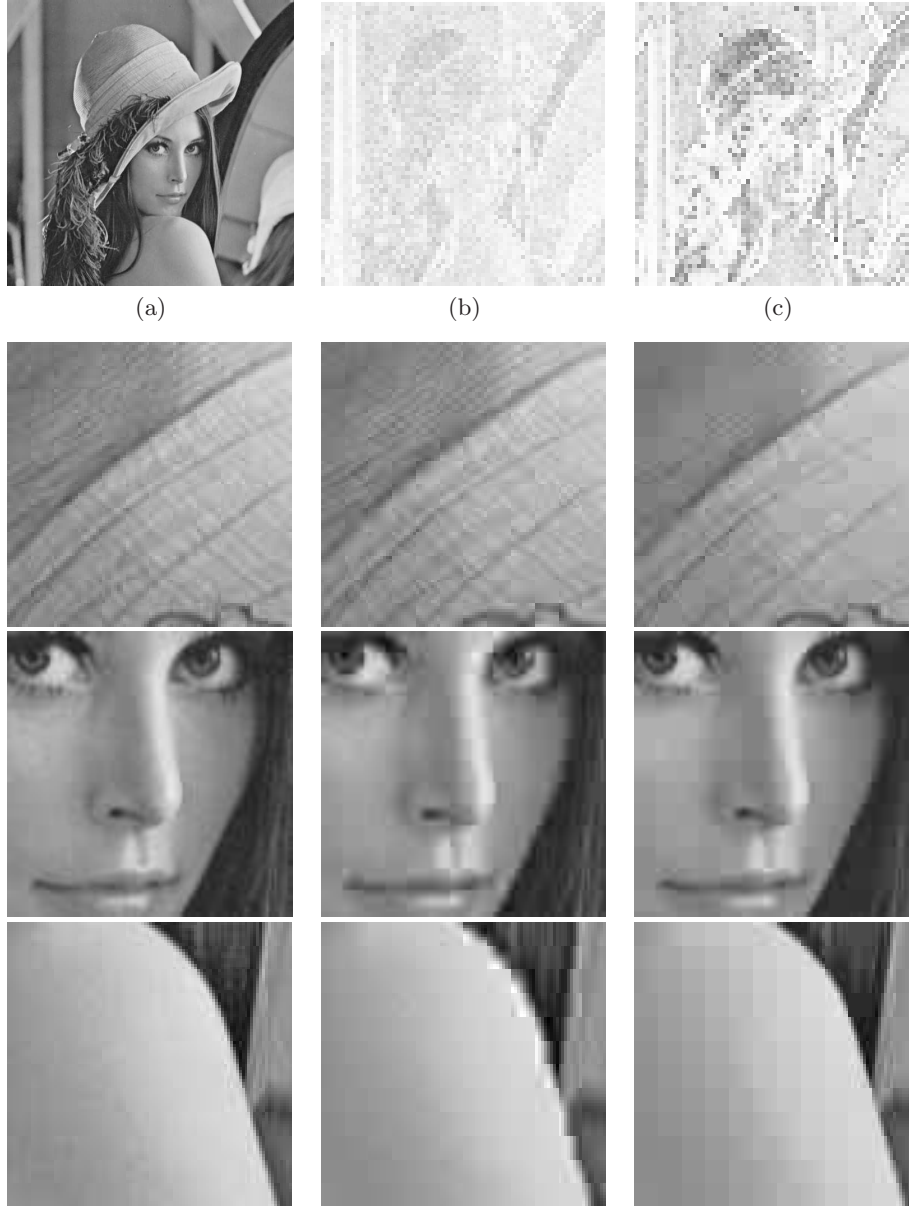


Fig. 1. Top row. (a): Original 512×512 8bpp *Lena* image. (b): BSSIM map of BSSIM-optimized block-DCT approximation of *Lena* with 10,000 non-zero higher-order DCT coefficients. (c): BSSIM map of L^2 -optimized block DCT approximation of *Lena* with 10,000 non-zero higher-order DCT coefficients. **Bottom rows.** First column: Patches of the original image. Second column: Patches of the BSSIM-optimized image. Third column: Patches of the L^2 -optimized image. Second row: Hat, cropped from (81, 169) to (176, 264). Third row: Face, cropped from (257, 257) to (362, 362). Fourth row: Shoulder, cropped from (401, 289) to (496, 384).

sponding to only one higher-order coefficient), it is not too hard to show that, for small ϵ_2 and nonzero $s_{\mathbf{x}}$, $\alpha \approx s_{\mathbf{x}}$. This implies that blocks with large variance – typically those with edges – can have large α values when a low number of coefficients is being used. This feature is demonstrated in the figure, particularly along the shoulder.

In Fig. 2 are presented plots of the BSSIM values vs. the number of non-zero coefficients for the SSIM- and L^2 -based optimization methods. For this particular example, we see that the greatest increase in BSSIM from the L^2 -based method occurs when the number K of non-zero higher-order coefficients is between 2000 and 3000. As K increases toward 10,000, the difference between the two methods decreases.

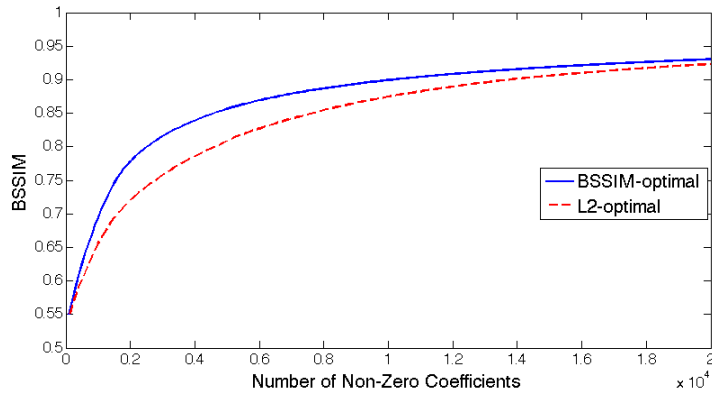


Fig. 2. Graphs of the BSSIM values of the BSSIM-optimized and L^2 -optimized 8×8 block-DCT approximations of *Lena* image, as a function of the number of non-zero higher-order DCT coefficients employed.

4 Concluding remarks

We have mathematically examined the problem of SSIM-based approximations of signal/image blocks in terms of an orthogonal basis. In general, the non-convex nature of the SSIM function complicates optimization problems. In this case, however, a closed-form solution can be found since the partial derivatives of the SSIM function with respect to the expansion coefficients are rather straightforward and stationary points can be determined. The optimal SSIM approximation is found to be related to the well-known optimal L^2 approximation in a quite simple manner: The zeroth-order coefficients of both expansions are the same, and higher-order SSIM coefficients are obtained from their Fourier counterparts by simple scaling. On its own, this algorithm does not yield any significant improvement in the usual L^2 -optimized expansions. But we have shown its potential use

in the construction of approximations that are subject to fixed “bit budgets.” Only a simple example was studied here – there are many possibilities for further exploration, including coding/compression and enhancement.

The results of Section 2 have inspired an investigation of more general cases which will be reported elsewhere. These cases include a generalization of the basis set to the nonorthonormal and overcomplete cases. Other avenues of future research include the combination of SSIM-optimization with prior image models in a Bayesian framework.

Acknowledgements This research was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC) in the form of a Postgraduate Scholarship (DB), Discovery Grants (ERV, ZW) and a Strategic Grant (ZW). It was also supported in part by the Ontario Ministry of Research & Innovation in the form of an Early Researcher Award (ZW).

References

1. S. S. Channappayya, A. C. Bovik, C. Caramanis, and R. W. Heath. Design of linear equalizers optimized for the structural similarity index. *IEEE Trans. Image Processing*, 17(6):1624–1639, June 2008.
2. S. S. Channappayya, A. C. Bovik, C. Caramanis, and R. W. Heath. SSIM-optimal linear image restoration. In *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing*, Las Vegas, Nevada, March 2008.
3. S. S. Channappayya, A. C. Bovik, and R. W. Heath. Rate bounds on ssim index of quantized images. *IEEE Trans. Image Processing*, 17(6):857–872, September 2008.
4. N. Ponomarenko and K. Egiazarian. Tampere Image Database 2008 TID2008. <http://www.ponomarenko.info/tid2008.htm>.
5. T. Richter and K. J. Kim. A MS-SSIM optimal JPEG 2000 encoder. In *Data Compression Conference*, pages 401–410, Snowbird, Utah, March 2009.
6. H. R. Sheikh, K. Seshadrinathan, A. K. Moorthy, Z. Wang, A. C. Bovik, and L. K. Cormack. Image and video quality assessment research at LIVE. <http://live.ece.utexas.edu/research/quality/>.
7. M. J. Wainwright, O. Schwartz, and E. P. Simoncelli. Natural image statistics and divisive normalization: Modeling nonlinearity and adaptation in cortical neurons. In R. Rao, B. Olshausen, and M. Lewicki, editors, *Probabilistic Models of the Brain: Perception and Neural Function*, pages 203–222. MIT Press, 2002.
8. Z. Wang and A. C. Bovik. Mean squared error: Love it or leave it? A new look at signal fidelity measures. *IEEE Signal Processing Magazine*, 26(1):98–117, January 2009.
9. Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Trans. Image Processing*, 13(4):600–612, April 2004.
10. Z. Wang, Q. Li, and X. Shang. Perceptual image coding based on a maximum of minimal structural similarity criterion. In *Proc. IEEE Int. Conf. Image Proc.*, San Antonio, TX, September 2007.
11. Z. Wang and E. P. Simoncelli. Maximum differentiation (MAD) competition: A methodology for comparing computational models of perceptual quantities. *Journal of Vision*, 8(12):1–13, 2008.