Chapter in Handbook of Convex Optimization Methods in Imaging Science, Springer International Publishing, 2017

Optimizing Image Quality

Dominique Brunet, Sumohana S. Channappayya, Zhou Wang, Edward R. Vrscay and Alan C. Bovik

Abstract The fact that multimedia services have become *the* major driver for next generation wireless networks underscores their technological and economic impact. A vast majority of these multimedia services are consumer-centric and therefore must guarantee a certain level of perceptual quality. Given the massive volumes of image and video data in question, it is only natural to adopt automatic quality prediction and optimization tools. The past decade has seen the invention of several excellent automatic quality prediction tools for natural images and videos. While these tools predict perceptual quality scores accurately, they do not necessarily lend themselves to standard optimization techniques. In this chapter, a systematic framework for optimization with respect to a perceptual quality assessment algorithm is presented. The Structural SIMilarity (SSIM) index, which has found vast commercial acceptance owing to its high performance and low complexity, is the representative image quality assessment model that is studied. Specifically, a detailed exposition of the mathematical properties of the SSIM index is presented first, followed by a discussion on the design of linear and non-linear SSIM-optimal image restoration algorithms.

Dominique Brunet
University of Waterloo, e-mail: dbrunet@uwaterloo.ca

Sumohana S. Channappayya
Indian Institute of Technology Hyderabad, e-mail: sumohana@iith.ac.in

Zhou Wang
University of Waterloo, e-mail: Z.Wang@ece.uwaterloo.ca

Edward R. Vrscay
University of Waterloo, e-mail: ervrscay@uwaterloo.ca

Alan C. Bovik
The University of Texas at Austin, e-mail: bovik@ece.utexas.edu

1 Introduction

Optimization of perceptual image quality has traditionally been associated with image compression and is aimed at achieving the highest perceptual quality at the lowest possible encoding rate. The earliest known design of perceptually optimal algorithms can be traced to Mannos and Sakrison's seminal work [21] on image coding with respect to a visual fidelity criterion. The primary challenge with perceptual optimization is the fact that a majority of the state-of-the-art perceptual quality algorithms do not enjoy convenient mathematical properties such as differentiability, convexity and metricity. Further, several of these algorithms employ parameters and thresholds that are again not easy to deal with in an optimization framework. While these challenges appear to make the problem intractable, we present a systematic framework in this chapter to address the problem of perceptual optimization. As a particular and particularly pratical example to illustrate the framework, we work with the Structural SIMilarity (SSIM) index.

1.1 Image Quality Assessment Measures

The goal of *image quality assessment* (IQA) is to predict the quality of images in a manner that is consistent with human subjective evaluation. IQA can be divided into *full-reference*, *reduced-reference* and *no-reference*, depending on the full, reduced and non-availability of the original (ground truth) image. Full-reference IQA algorithms are more of image fidelity predictors since the goal is to measure similarity between two images. They are often used as quality predictors when one of the images being compared is considered to have pristine quality. By contrast, reduced- and no-reference IQA algorithms rely more on prior knowledge about high quality natural images in the statistical sense [36].

The focus in this chapter will be limited to full-reference algorithms, and specifically to those methods that are based on the notion of *structural similarity*. A brief chronological evolution of the Structural SIMilarity (SSIM) index is presented next.

The Universal Image Quality Index (UIQI) [35] was the precursor to the successful SSIM index [37] and introduced the idea of measuring local luminance, contrast and structural similarity between a reference image and its test version. For an original image signal $\mathbf{x} = \{x_i \mid i = 1,\ldots,N\}$ and its distorted version $\mathbf{y} = \{y_i \mid i = 1,\ldots,N\}$, UIQI is defined as

$$Q(\mathbf{x},\mathbf{y}) = \left(\frac{2\mu_x\mu_y}{\mu_x^2+\mu_y^2}\right)\left(\frac{2\sigma_x\sigma_y}{\sigma_x^2+\sigma_y^2}\right)\left(\frac{\sigma_{xy}}{\sigma_x\sigma_y}\right), \tag{1}$$

where μ_x is the mean of the image signal, σ_x^2 its variance and σ_{xy} the covariance between the original and distorted version of the image signal. The three terms in this formulation are the foundations of structural similarity based image quality assess-

ment. The first term measures luminance similarity between the images, the second measures contrast similarity and the third term measures correlation or *structural similarity* between the images. If the quality metric were to be applied on a patch-wise basis to the images, the overall image quality is estimated as

$$Q = \frac{1}{M}\sum_{j=1}^{M} Q_j, \tag{2}$$

where Q_j is the UIQI of the j^{th} image patch and M is the number of patches.

The Structural SIMilarity (SSIM) index builds on the ideas introduced by UIQI by using weighted means, variances and covariances in addition to introducing stabilizing constants to Q so as to avoid numerical problems when the local statistics are close to zero. The local SSIM index is defined as [37]

$$\text{SSIM}(\mathbf{x},\mathbf{y}) = \left(\frac{2\mu_x\mu_y + C_1}{\mu_x^2 + \mu_y^2 + C_1}\right)\left(\frac{2\sigma_x\sigma_y + C_2}{\sigma_x^2 + \sigma_y^2 + C_2}\right)\left(\frac{\sigma_{xy} + C_3}{\sigma_x\sigma_y + C_3}\right), \tag{3}$$

where C_1, C_2 and C_3 are stabilizing constants. Also, the mean, variance and covariance are estimated locally via

$$\mu_x = \mathbf{w}^T\mathbf{x}, \tag{4}$$
$$\sigma_x^2 = (\mathbf{x} - \mu_x\mathbf{e})^T \text{diag}(\mathbf{w})(\mathbf{x} - \mu_x\mathbf{e}), \tag{5}$$
$$\sigma_{xy} = (\mathbf{x} - \mu_x\mathbf{e})^T \text{diag}(\mathbf{w})(\mathbf{y} - \mu_y\mathbf{e}), \tag{6}$$

where $\mathbf{w}$ is a normalized weight vector and $\mathbf{e}$ is a vector of ones.

The SSIM index is also expressed as

$$\text{SSIM}(\mathbf{x},\mathbf{y}) = l(\mathbf{x},\mathbf{y})c(\mathbf{x},\mathbf{y})s(\mathbf{x},\mathbf{y}), \tag{7}$$

where $l(\mathbf{x},\mathbf{y})$ is the luminance term, $c(\mathbf{x},\mathbf{y})$ is the contrast term and $s(\mathbf{x},\mathbf{y})$ corresponds to the structure term. As shown in [37], SSIM embodies important masking mechanisms in the terms $l(\mathbf{x},\mathbf{y})$ and $c(\mathbf{x},\mathbf{y})$, specifically luminance masking (Weber's law) and contrast masking, both of which are key determinants of image quality. The structure term $s(\mathbf{x},\mathbf{y})$ captures the notion that image distortion destroys perceptually relevant image structure.

In its most general form, the SSIM index is defined as

$$\text{SSIM}(\mathbf{x},\mathbf{y}) = [l(\mathbf{x},\mathbf{y})]^{\alpha}[c(\mathbf{x},\mathbf{y})]^{\beta}[s(\mathbf{x},\mathbf{y})]^{\gamma}, \tag{8}$$

where the positive exponents α, β, γ determine the importance assigned to each of the components. Since the structure term can be negative, γ should be normalized to 1 to avoid any complex number and not to favorize any anti-correlation. Alternatively, $s(\mathbf{x},\mathbf{y})$ may be transformed to $\max(0, s(\mathbf{x},\mathbf{y}))$ to avoid any negative values.

The SSIM index is applied block-wise to arrive at the image level SSIM index. Again, as with UIQI,

$$\mathrm{SSIM}(\mathbf{x},\mathbf{y}) = \frac{1}{M}\sum_{j=1}^{M}\mathrm{SSIM}_j(\mathbf{x},\mathbf{y}), \tag{9}$$

where SSIM_j is the SSIM index of the j^{th} image patch and M is the number of patches.

The Multi-Scale-SSIM (MS-SSIM) index [41] is an improvement over the SSIM index in that it measures structural similarity over multiple spatial scales. The MS-SSIM index is defined as

$$\mathrm{SSIM}(\mathbf{x},\mathbf{y}) = [l_J(\mathbf{x},\mathbf{y})]^{\alpha_J}\prod_{j=1}^{J}[c_j(\mathbf{x},\mathbf{y})]^{\beta_j}[s_j(\mathbf{x},\mathbf{y})]^{\gamma_j} \tag{10}$$

corresponding to J spatial scales. Starting from the highest spatial scale corresponding to the original image resolution, successive spatial scales (at lower resolution) are obtained by decimation i.e., low-pass filtering of the current scale followed by downsampling by a factor of 2. Following the notation established earlier, $l_J(\mathbf{x},\mathbf{y})$ corresponds to luminance similarity after $(J-1)$ stages of decimation. In similar fashion, $c_j(\mathbf{x},\mathbf{y})$ and $s_j(\mathbf{x},\mathbf{y})$ correspond to contrast and structural similarity after $j-1$ decimation stages respectively. Further, α_J is the exponent applied to the luminance term while β_j and γ_j correspond to exponents applied to the contrast and structure terms at the $(j-1)^{\text{st}}$ decimation stage, respectively.

Many extensions of the SSIM index have been proposed in the literature. To name a few, we mention the Complex-Wavelet-SSIM (CW-SSIM) [34], Information content Weighted SSIM (IW-SSIM) [38], SSIM extension for color images [22, 19] and for videos [39].

A comprehensive comparison of IQA measures was performed in the TID-2008 experiment [29]. MS-SSIM was the clear winner, followed by SSIM at the second place. Several other metrics with better performance on this database have been introduced since then, many of them inspired by SSIM. According to the TID-2013 experiment [28], Feature SIMilarity (FSIM) [43], Sparse Feature Fidelity (SFF) [8], PNSR-HA/HMA [27], Block-Based Multi-Metric (BMMF) [17] and Spectral Residual SIMilarity (SR-SIM) [42] correlate better with subjective quality assessment than MS-SSIM on the TID database.

Yet, SSIM and MS-SSIM have found vast commercial and academic acceptance, owing to its high efficiency in terms of computation (it is very fast) and performance (it correlates with human judgments nearly as well as the top models). Indeed, SSIM and MS-SSIM are marketed and used throughout the global broadcast, cable and satellite television industries. The SSIM team members each received a Primetime Emmy Award in October 2015 for their work. Morever, because of the simplicity of the SSIM index, it is a good prototype of a perceptual optimization criterion. It is hoped that a detailed study of optimization techniques for the SSIM index will inspire similar studies for future IQA models and for the design of improved IQA methods that will better serve the multimedia industry.

1.2 Perceptual optimization framework

Perceptual optimization is a particular case of the general optimization framework where the objective measure models perceptual quality of an image. Before studying the specific case of SSIM-based optimization, we lay the groundwork for perceptual optimization in both Bayesian and variational perspectives.

We first consider the no-reference IQA case. Let $\mathbf{x}$ be an image and let $Q(\mathbf{x})$ be a quality measure of $\mathbf{x}$ with $Q(\mathbf{x}) \geq 0$ and $Q(\mathbf{x}) = 0$ for perfect quality. Given a distorted image $\mathbf{y}$ obtained from an unknown image $\mathbf{x}$ by a model of distortion $\mathbf{y} = D(\mathbf{x})$, we want to restore the original image $\mathbf{x}$. To do so, we seek an image $\hat{\mathbf{x}}$ that will optimize the quality criterion $Q(\hat{\mathbf{x}})$ with the constraint $\mathbf{y} = D(\hat{\mathbf{x}})$. Often, the distortion model will include a stochastic component such as additive noise. In that case, the distortion model can be expressed as $\mathbf{y} = D(\mathbf{x}) + \xi$, where D is the deterministic part and ξ is the noise component. For example, for additive white Gaussian noise, the probability distribution function of ξ will be

$$P(\Xi = \xi) \propto \exp(-1/2\|\xi\|_2^2/\sigma^2),$$

where σ^2 is the variance of the noise. We can also assume a probabilistic model for the unknown image $\mathbf{x}$ such as

$$P(X = \mathbf{x}) \propto \exp(-Q(\mathbf{x})^\alpha/\beta).$$

Note that other monotonic transformations of $Q(\mathbf{x})$ could be found instead of a power transformation. Since $\xi = \mathbf{y} - D(\mathbf{x})$ the probability distribution function of ξ can also be seen as the conditional probability of $\mathbf{y}$ given $\mathbf{x}$:

$$P(Y = \mathbf{y}|X = \mathbf{x}) \propto \exp(-1/2\|\mathbf{y} - D(\mathbf{x})\|_2^2/\sigma^2).$$

We can now use the Bayes' rule to find the probability distribution of the clean image $\mathbf{x}$ given the distorted image $\mathbf{y}$:

$$P(X = \mathbf{x}|Y = \mathbf{y}) \propto P(Y = \mathbf{y}|X = \mathbf{x})P(X = \mathbf{x}) \tag{11}$$

$$\propto \exp(-1/2\|\mathbf{y} - D(\mathbf{x})\|_2^2/\sigma^2)\exp(-Q(\mathbf{x})^\alpha/\beta). \tag{12}$$

The maximum *a posteriori* estimator is then obtained by finding the optimal image $\mathbf{x}$. Taking the negative of the logs, it yields to

$$-\log(P(X = \mathbf{x}|Y = \mathbf{y})) \propto 1/2\|\mathbf{y} - D(\mathbf{x})\|_2^2/\sigma^2 + Q(\mathbf{x})^\alpha/\beta.$$

This is the variational form for the image restoration problem. Notice that the mean squared error in the first term does not come from the image quality model, but rather the model of the noise. The second term, usually called the regularization term, often adopts simple forms such as the L^2-norm (for Tikhonov regularization), the total variation or the norm of the gradient or of second-order partial derivatives

(for thin-plate spline). From a Bayesian perspective, these represent different models of image quality.

For the full-reference case, we can use a distance (or dissimilarity) measure $d(\mathbf{x},\mathbf{z})$. The goal of the optimization will then be to bring the estimated clean image $\hat{\mathbf{x}}$ close to a prior image $\mathbf{z}$. For example, we could minimize $d(\mathbf{x},\mathbf{z})$ over all $\mathbf{x}$ satisfying the constraint $\mathbf{y} = D(\mathbf{x})$. Instead of a single prior image, we could also have a dictionary of images $\mathbf{z}_1, \mathbf{z}_2, \ldots, \mathbf{z}_P$. In this case, we will solve the problem

$$\hat{\mathbf{x}} = \arg\min_{p} \min_{\mathbf{x}:\mathbf{y}=D(\mathbf{x})} d(\mathbf{x},\mathbf{z}_p). \tag{13}$$

If instead we have a (possibly empirical) prior probability distribution $P(\mathbf{z})$, with a stochastic model of distortion $P(Y = \mathbf{y}|Z = \mathbf{z})$, then Bayes' formulas will lead to

$$P(Z = \mathbf{z}|Y = \mathbf{y}) \propto P(Y = \mathbf{y}|Z = \mathbf{z})P(Z = \mathbf{z}).$$

Finding the maximum *a posteriori* estimator does not require any optimization of an image quality measure. (Or we could say that the image quality measure was empirically found to be $Q(\mathbf{z})^{\alpha} \propto -\log(P(Z = \mathbf{z}))$.) However, we could instead look to maximize the expected perceptual quality:

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} E_Z[d(\mathbf{z},\mathbf{x})|Y = \mathbf{y}] \tag{14}$$

$$= \arg\min_{\mathbf{x}} \int d(\mathbf{z},\mathbf{x})P(Z = \mathbf{z}|Y = \mathbf{y})d\mathbf{z}. \tag{15}$$

The optimization problems presented in this chapter will roughly follow one of the forms presented above. However, by sampling through the SSIM-optimization literature, we will not present a completly methodical and exhaustive approach as it was outlined here.

1.3 Chapter Overview

We first review the mathematical properties of SSIM that makes it a suitable criterion for optimization. We then study several optimization problems starting with a local (block-based) perspective. Specifically, we discuss SSIM-optimal equalizer design, SSIM-optimal soft-thresholding algorithms, and SSIM best basis approximation. We then present a variational formula with a SSIM term that allows to pass from a local (block-based) to a global (image-wide) solution. We illustrate the techniques with several examples and compare the mean-square error equivalent.

2 Mathematical Properties of the SSIM index

2.1 Structural and non-structural distortions

The main insight behind the Structural Similarity index and the family of related quality measures is a decomposition of images/signals into structural and non-structural components. Indeed, it is generally observed that for the same mean-square error, structural image distortions are perceived more strongly than non-structural ones.

We shall refer to distortions that do not alter the general shape of an image as being "non-structural". Change in luminance, change in contrast, translation and rotation are some examples of non-structural distortions. We want perceptual metrics to be quasi-invariant to these types of distortions.

On the other hand, we shall refer to distortions that strongly affect the perceptual quality of an image as being "structural". As we will see, all remaining distortions will be lumped together once the non-structural distortions have been accounted for.

A commonly used simplification of the local SSIM (3) is obtained by setting $C_3 = C_2/2$. The formula then reduces to

$$\mathrm{SSIM}(\mathbf{x},\mathbf{y}) = \left(\frac{2\mu_x\mu_y + C_1}{\mu_x^2+\mu_y^2+C_1}\right)\left(\frac{2\sigma_{x,y}+C_2}{\sigma_x^2+\sigma_y^2+C_2}\right), \tag{16}$$

$$= S_1(\mathbf{x},\mathbf{y})S_2(\mathbf{x},\mathbf{y}). \tag{17}$$

In this simplified form of the SSIM index, only one type of non-structural distortion is considered explicitly: the luminance shift. Given a grayscale image patch $\mathbf{x}$, its luminance is a function of the average value μ_x of the patch. The decomposition of an image patch into structural and non-structural parts is then given by:

$$\mathbf{x} = \mu_x\mathbf{e} + (\mathbf{x} - \mu_x\mathbf{e}),, \tag{18}$$

where $\mathbf{e} = (1,1,\ldots,1)$ is in the direction of the mean and $\mathbf{x} - \mu_x\mathbf{e}$ is the zero-mean component. The luminance term S_1 of the SSIM index acts on the mean direction of the signal, whereas the combined contrast-correlation term S_2 involves the zero-mean component of the signal.

We can then rewrite the components S_1 and S_2 using the two following projections:

$$P_1(\mathbf{x}) = \mu_x\mathbf{e}, \text{ and} \tag{19}$$

$$P_2(\mathbf{x}) = (\mathrm{Id} - P_1)\mathbf{x} = \mathbf{x} - \mu_x\mathbf{e}, \tag{20}$$

where $\mathrm{Id}(\mathbf{x}) = \mathbf{x}$ is the identity operator (matrix). For $\mathbf{x},\mathbf{y} \in \mathbf{R}^N$, we define

$$d(\mathbf{x},\mathbf{y}) = \left(\frac{\|\mathbf{x}-\mathbf{y}\|_2^2}{\|\mathbf{x}\|_2^2+\|\mathbf{y}\|_2^2+C}\right)^{1/2}$$

for some positive constant C. This is a valid normalized distance metric [7]. Note that

$$1-S_1(\mathbf{x},\mathbf{y}) = \frac{|\mu_x-\mu_y|^2}{\mu_x^2+\mu_y^2+C_1} = [d(P_1\mathbf{x},P_1\mathbf{y})]^2 \text{ and} \tag{21}$$

$$1-S_2(\mathbf{x},\mathbf{y}) = \frac{\|(\mathbf{x}-\mu_x\mathbf{e})-(\mathbf{y}-\mu_y\mathbf{e})\|_2^2}{\|\mathbf{x}-\mu_x\mathbf{e}\|_2^2+\|\mathbf{y}-\mu_y\mathbf{e}\|_2^2+(N-1)C_2} = [d(P_2\mathbf{x},P_2\mathbf{y})]^2. \tag{22}$$

This implies that $1-S_2(\mathbf{x},\mathbf{y})$ can be interpreted as an inverse variance weighted normalized mean square error. The fundamental difference between $1-S_2(\mathbf{x},\mathbf{y})$ and the normalized mean square error $\|P_2\mathbf{x}-P_2\mathbf{y}\|_2^2$ is that the former models the masking effect by weighting more heavily distortions on flatter patches.

By convention, optimization problems are cast as minimization problems. This is of course always possible to pass from maximization to minimization since $\max_{\mathbf{x}} f(\mathbf{x}) = \min_{\mathbf{x}} af(\mathbf{x})+b$ with $a<0$. We thus consider the minimization of $f_1(\mathbf{x},\mathbf{y}) := 1-S_1(\mathbf{x},\mathbf{y})$ and $f_2(\mathbf{x},\mathbf{y}) := 1-S_2(\mathbf{x},\mathbf{y})$.

2.2 Convexity and quasi-convexity

The functions f_1 and f_2 that were derived from the components of the simplified SSIM are not convex everywhere. As it was studied in detail in [7], for fixed $\mathbf{y}$, f_1 is convex for $\{\mathbf{x} : 0 \le P_1\mathbf{x} \le \sqrt{3}P_1\mathbf{y}\}$ and f_2 is convex for all points in $\{\mathbf{x} : \|P_2\mathbf{x}-P_2\mathbf{y}\|^2 \le \|P_2\mathbf{y}\|^2(\sqrt{3}-1)^2\}$. The exact limit of the region of convexity is actually a tear-shaped region which is rotated around the direction of $P_2\mathbf{y}$.

Although the functions f_1 and f_2 are not convex everywhere, they possess a weaker form of convexity called quasi-convexity.

Definition 1. Given a convex set X, a function $f : X \to \mathbf{R}$ is said to be *quasi-convex* if its *h-sublevel set*, defined as

$$X_h = \{x \in X | f(x) \le h\}, \tag{23}$$

is a convex set for all $h \in Range(f)$.

Any convex function is necessarily quasi-convex. An example of function that is quasi-convex but not convex is $f(x) = |x|^{1/2}$.

Quasi-convexity is a useful property for non-linear optimization [4]:

Theorem 2.1 *Let X be a convex vector space and let $f : X \to \mathbf{R}$ be a quasi-convex function. If f has a minimum, then it is either unique or the function is constant in a neighborhood of the minimum.*

It thus means that a function f will have a unique minimum if it is monotonically increasing away from the minimum.

The quasi-convexity of f_1 for $P_1\mathbf{x} \ge 0, P_2\mathbf{y} \ge 0$ and of f_2 on the half-plane $(P_2\mathbf{y})^T P_2\mathbf{x} \ge 0$ is easy to prove. The set of points such that $f_i \le h < 1$ for $i = 1,2$

is, respectively, the interval (hyperball) centered at $\frac{P_i\mathbf{y}}{1-h^2}$ of mid-length (radius) of $\sqrt{\|P_i\mathbf{y}\|^2\left(\frac{1}{(1-h^2)^2}-1\right)+C_i\frac{h^2}{1-h^2}}$, which is a convex set.

2.3 Combination rule

In order to optimize perceptual quality, we need either to optimize the vector-valued function (f_1, f_2) or to devise some way to collapse the vector to a scalar function. We then examine each option in terms of the (quasi-)convexity of the resulting function.

In the simplified form of SSIM, the product between S_1 and S_2 is taken. Other combination rules could also be devised, and the choice we make will affect not only how good our perceptual image quality model is, but also the mathematical properties of the created function.

Writing the local SSIM in term of $f_1 = 1 - S_1$ and $f_2 = 1 - S_2$, we get

$$\begin{aligned} \mathrm{SSIM}(\mathbf{x},\mathbf{y}) &= (1 - f_1(\mathbf{x},\mathbf{y}))(1 - f_2(\mathbf{x},\mathbf{y})) && (24)\\ &= 1 - f_1(\mathbf{x},\mathbf{y}) - f_2(\mathbf{x},\mathbf{y}) + f_1(\mathbf{x},\mathbf{y})f_2(\mathbf{x},\mathbf{y}). && (25)\end{aligned}$$

Maximizing SSIM is equivalent to minimizing the following function,

$$1 - \mathrm{SSIM}(\mathbf{x},\mathbf{y}) = f_1(\mathbf{x},\mathbf{y}) + f_2(\mathbf{x},\mathbf{y}) - f_1(\mathbf{x},\mathbf{y})f_2(\mathbf{x},\mathbf{y}). \quad (26)$$

Another simple combination rule would be to take the sum of f_1 and f_2:

$$F_1(\mathbf{x},\mathbf{y}) = f_1(\mathbf{x},\mathbf{y}) + f_2(\mathbf{x},\mathbf{y}) \quad (27)$$

This can be seen as a linear approximation of 1-SSIM. Surprisingly, this approximation performed slightly better in a psycho-visual experiment [5]. Moreover, this alternative form preserves the property of a metric [7].

If $\mathbf{x}$ is in the region of convexity for f_1 and f_2, then it will be automatically in the region of convexity of F_1. However, contrary to the convex function case, the sum of quasi-convex functions is not necessarily quasi-convex. So F_1 does not inherit the quasi-convexity property of f_1 and f_2.

Even if the scalarized function to be optimized is neither convex nor quasi-convex, it is still sometimes possible to optimize each component independently. Indeed, the decomposition given in (18) is an orthogonal decomposition. The Orthogonal Decomposition Theorem goes as follows:

Theorem 1. *Let X be a space and $\{P_k\}_{k=1}^{K-1}$ be orthogonal projections. Define $P_K = Id - \sum_{k=1}^{K-1} P_k$, where $Id(x) = x$ is the identity function. Then each point $x \in X$ has a unique decomposition as*

$$\mathbf{x} = \sum_{k=1}^{K} P_k(\mathbf{x}).$$

In particular,

$$\mathbf{x} = \mathbf{y} \iff P_k(\mathbf{x}) = P_k(\mathbf{y}) \text{ for } 1 \leq k \leq K.$$

For the particular case of SSIM, the Orthogonal Decomposition Theorem says that if $\mu_x = \mu_y$ and $\mathbf{x} - \mu_x e = \mathbf{y} - \mu_y e$, then $\mathbf{x} = \mathbf{y}$, which is quite trivial. However, this kind of result will also be valid for more complex types of orthogonal decomposition. The next proposition follows immediately from the Orthogonal Decomposition Theorem:

Theorem 2. *Let $X = X_1 + X_2$ be an orthogonal decomposition of X and let f_1 on X_1 and f_2 on X_2 be two functions. Then*

$$\min_{\mathbf{z}_1 : \mathbf{z}_1 = P_1 \mathbf{z}_1} f_1(\mathbf{z}, P_1\mathbf{y}) + \min_{\mathbf{z}_2 : \mathbf{z}_2 = P_2 \mathbf{z}_2} f_2(\mathbf{z}_2, P_2\mathbf{y}) = \min_{\mathbf{z}} f_1(P_1\mathbf{z}, P_1\mathbf{y}) + f_2(P_2\mathbf{z}, P_2\mathbf{y}). \quad (28)$$

Moreover, if $f_1, f_2 > 0$, then

$$\min_{\mathbf{z}_1 : \mathbf{z}_1 = P_1 \mathbf{z}_1} f_1(\mathbf{z}, P_1\mathbf{y}) \min_{\mathbf{z}_2 : \mathbf{z}_2 = P_2 \mathbf{z}_2} f_2(\mathbf{z}_2, P_2\mathbf{y}) = \min_{\mathbf{z}} f_1(P_1\mathbf{z}, P_1\mathbf{y}) f_2(P_2\mathbf{z}, P_2\mathbf{y}). \quad (29)$$

This theorem can be also used for a subdomain $A \subset X$, but with the limitation that $A = P_1 A + P_2 A$.

2.4 Spatial aggregation

The simplest way to spatially aggregate local SSIM scores is by averaging (9). Other pooling options could be possible such as a power mean or maximum error [40]. Visual saliency models can also be used to guide the pooling strategy [44]. Finally, multi-resolution schemes such as MS-SSIM can be devised.

In order to preserve the quasi-convexity property of the components of SSIM, an interesting option would be to take the maximum value of the scores on local patches:

$$\text{SSIM}_{\max} = \max_{1 \leq j \leq M} \text{SSIM}_j(\mathbf{x}, \mathbf{y}). \quad (30)$$

This strategy can be justified empirically by the fact that the eye will be attracted to the most salient feature, which could be said to be the one with the greatest dissimilarity. The advantage of this choice is that the aggregated score will be quasi-convex, since it is the maximum of quasi-convex functions.

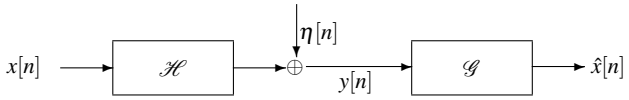

Fig. 1 Block diagram of a general linear time invariant equalizer system. The goal is to design a linear equalizer block $\mathscr{G}$ that maximizes the SSIM index between the source process $x[n]$ and the restored process $\hat{x}[n]$ is maximized. To solve this problem it is assumed that the LTI filter $\mathscr{H}$, and the power spectral density of the noise process $\eta[n]$ are known.

3 Perceptually Optimal Algorithm Design

3.1 SSIM-optimal Equalizer Design

One of the oldest and most widely researched problems in digital image processing is image restoration, a.k.a. equalization [1], which has its origins in the early days of America's space program. Given the problem's rich history, there exist several excellent solutions that form a part of most image acquisition systems [18]. In this section, we discuss a relatively recent image restoration solution that explicitly optimizes the SSIM index. A precursor to this solution is the SSIM-optimal linear estimator that addresses the image denoising problem [10].

3.2 Equalization Problem

The equalization problem is as follows: Design an equalizer $g[n]$ of length N (for any N) that optimizes the SSIM index between the reference and restored wide sense stationary (WSS) processes $x[n]$ and $\hat{x}[n]$ respectively, given the observed process $y[n]$ that is a blurred and noisy version of $x[n]$. The blurring filter $h[n]$ of length M is assumed to be linear and time-invariant (LTI) and the noise process $\eta[n]$ is assumed to be additive and white in nature. Furthermore, it is assumed that the blurring filter $h[n]$ and the power spectral density (PSD) of $\eta[n]$ are known at the receiver. This system is summarized in Fig. 1 and the optimization problem is set up as follows. Given

$$y[n] = h[n] * x[n] + \eta[n], \tag{31}$$

design a filter $g[n]$ of length N such that

$$g^*[n] = \arg \max_{g[n] \in \mathbb{R}^N} \text{SSIM}(x[n], \hat{x}[n]), \tag{32}$$

where

$$\hat{x}[n] = g[n] * y[n]. \tag{33}$$

3.3 Solution

For completeness, the standard SSIM index [37] as defined for deterministic signals is reviewed first. However, the equalizer design problem in (32) is defined on WSS random processes thereby rendering a direct application of the deterministic SSIM index infeasible. To circumvent this issue, the definition of the SSIM index is extended to handle WSS processes. The optimization problem is then restated in terms of the extended definition of the SSIM index.

3.3.1 Equalization Problem Redefined

As noted previously, the definition of the SSIM index in (3) needs to be modified to measure similarity between WSS processes. This is accomplished via the definition of the *statistical SSIM index* (StatSSIM index).

Definition 2. Given two WSS random processes $x[n]$ and $y[n]$ with means μ_x and μ_y respectively, the statistical SSIM index is defined as

$$\begin{aligned}\text{StatSSIM}(x[n],y[n]) = &\left(\frac{2E[x[n]]E[y[n]]+C_1}{E[x[n]]^2+E[y[n]]^2+C_1}\right)\\ &\times\left(\frac{2E[(x[n]-E[x[n]])(y[n]-E[y[n]])]+C_2}{E[(x[n]-E[x[n]])^2]+E[(y[n]-E[y[n]])^2]+C_2}\right),\end{aligned} \tag{34}$$

where $E[]$ is the expectation operator. This is a straightforward extension of the pixel domain definition of the SSIM index by replacing sample means and variances with their statistical equivalents. This could be seen as a plug-in estimator of the expectation of SSIM (14). The problem in (32) is redefined as

$$g^*[n] = \arg\max_{g[n]\in\mathbb{R}^N} \text{StatSSIM}(x[n],\hat{x}[n]), \tag{35}$$

To solve (35), the StatSSIM index is first expressed in terms of the equalizer filter coefficients $g[n]$ using the definition in (34).

$$\begin{aligned}StatSSIM(x[n],\hat{x}[n]) = f(\mathbf{g}) &= \left(\frac{2\mu_x\mu_{\hat{x}}+C_1}{\mu_x^2+\mu_{\hat{x}}^2+C_1}\right)\left(\frac{2E[(x[n]-\mu_x)(\hat{x}[n]-\mu_{\hat{x}})]+C_2}{E[(x[n]-\mu_x)^2]+E[(\hat{x}[n]-\mu_{\hat{x}})^2]+C_2}\right)\\ &= \left(\frac{2\mu_x E[\sum_{i=0}^{N-1} g[i]y[n-i]]+C_1}{\mu_x^2+(E[\sum_{i=0}^{N-1} g[i]y[n-i]])^2+C_1}\right)\\ &\quad\times\left(\frac{2E[(x[n]-\mu_x)(\sum_{i=0}^{N-1} g[i]y[n-i]-E[\sum_{i=0}^{N-1} g[i]y[n-i]])]+C_2}{E[(x[n]-\mu_x)^2]+E[(\sum_{i=0}^{N-1} g[i]y[n-i]-E[\sum_{i=0}^{N-1} g[i]y[n-i]])^2]+C_2}\right)\end{aligned}$$

(which follows from the definition of convolution)

$$
\begin{aligned}
&= \left(\frac{2\mu_x \sum_{i=0}^{N-1} g[i]\mu_y + C_1}{\mu_x^2 + (\sum_{i=0}^{N-1} g[i]\mu_y)^2 + C_1} \right) \\
&\left(\frac{2E[(x[n]-\mu_x)(\sum_{i=0}^{N-1} g[i]y[n-i] - \sum_{i=0}^{N-1} g[i]\mu_y)] + C_2}{E[(x[n]-\mu_x)^2] + E[(\sum_{i=0}^{N-1} g[i]y[n-i] - \sum_{i=0}^{N-1} g[i]\mu_y)^2] + C_2} \right) \\
&\text{(since } x[n] \text{ is WSS, } h[n] \text{ is LTI, } y[n] \text{ is also WSS)} \\
&= \left(\frac{2\mu_x \mathbf{g}^T \mathbf{e}\mu_y + C_1}{\mu_x^2 + \mathbf{g}^T \mathbf{e}\mathbf{e}^T \mathbf{g}\mu_y^2 + C_1} \right) \\
&\left(\frac{2E[(x[n]-\mu_x)(\sum_{i=0}^{N-1} g[i](y[n-i]-\mu_y))] + C_2}{E[(x[n]-\mu_x)^2] + E[(\sum_{i=0}^{N-1} g[i](y[n-i]-\mu_y))^2] + C_2} \right) \\
&= \left(\frac{2\mu_x \mathbf{g}^T \mathbf{e}\mu_y + C_1}{\mu_x^2 + \mathbf{g}^T \mathbf{e}\mathbf{e}^T \mathbf{g}\mu_y^2 + C_1} \right) \left(\frac{2\mathbf{g}^T \mathbf{c}_{xy} + C_2}{\sigma_x^2 + \mathbf{g}^T \mathbf{K}_{yy} \mathbf{g} + C_2} \right)
\end{aligned}
\tag{36}
$$

where $\mathbf{g} = [g[0], g[1], \ldots, g[N-1]]^T, \mathbf{e} = [1, 1, \ldots, 1]^T$ are both length N vectors, μ_x, μ_y are the means of the source and observed processes respectively, $\mathbf{c}_{xy} = E[(x[n]-\mu_x)(\mathbf{y} - \mathbf{e}\mu_y)]$, is the cross-covariance between the source ($x[n]$) and the observed processes ($\mathbf{y} = (y[n], y[n-1], \ldots, y[n-(N-1)])^T$), σ_x^2 is the variance of the source process at zero delay, $\mathbf{K}_{yy} = E[(\mathbf{y} - \mathbf{e}\mu_y)(\mathbf{y} - \mathbf{e}\mu_y)^T]$, is the covariance matrix of size $N \times N$ of the observed process $y[n]$, and C_1, C_2 are stabilizing constants.

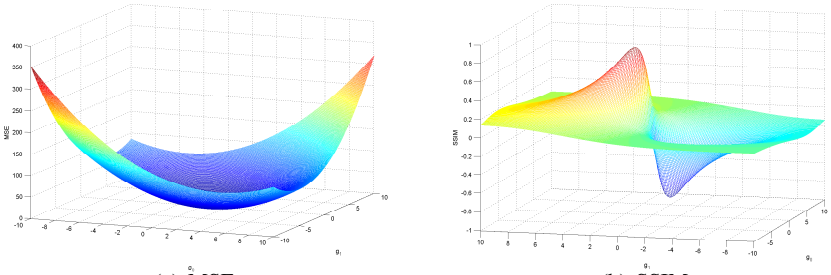

(a) MSE (b) SSIM

Fig. 2 MSE and StatSSIM index as a function of equalizer taps $g[0]$, $g[1]$.

From (36), it can be seen that the StatSSIM index is the ratio of a second degree polynomial to a fourth degree polynomial in g. It is now demonstrated that problem (36) admits a tractable, and in particular, near closed-form solution, with a complexity that is comparable to that of the minimum MSE solution.

3.3.2 StatSSIM-optimal linear equalization

The StatSSIM index is a non-convex function of $\mathbf{g}$ and local optimality conditions such as Karush-Kuhn-Tucker (KKT) cannot guarantee global optimality. The non-convex nature of the StatSSIM index is demonstrated in Fig. 2 and contrasted with the convex nature of MSE. In particular, any approach based on descent-type algorithms are likely to get stuck in local optima. To address this issue, the problem is transformed from its non-convex form into a quasi-convex formulation. Convex optimization problems are efficiently solvable using widely available optimization techniques and software [2, 4]. Moreover, we show below that a near-closed form solution can be achieved. In particular, the N-tap filter optimization is transformed into an optimization problem over only *two* variables for *any* N. Exploiting convexity properties, one can quickly search over one parameter by means of a bisection technique, thus reducing the problem to a univariate optimization problem. This last step can be efficiently performed using an analytic solution of a simplified problem.

3.3.3 Problem Reformulation

The problem is reformulated by noting that the first term of (36) (corresponding to the mean) is a function only of the sum of the filter coefficients $\mathbf{g}^T\mathbf{e}$. A typical constraint in filter design problems is that the filter coefficients add up to unity. The optimization problem in (36) is simplified by constraining $\mathbf{g}^T\mathbf{e} = \alpha$ and takes the form

$$\left[\begin{array}{l} \mathbf{g}(\alpha) = \text{argmax}_{\mathbf{g}\in\mathbf{R}^N}\left(\frac{2\mathbf{g}^T\mathbf{c}_{xy}+C_2}{\sigma_x^2+\mathbf{g}^T\mathbf{K}_{yy}\mathbf{g}+C_2}\right) \\ \text{subject to}: \mathbf{g}^T\mathbf{e} = \alpha. \end{array}\right], \tag{37}$$

The solution is now a function of α. The problem now changes to finding the highest StatSSIM index by searching over a range of α (typically in the interval $[1-\delta, 1+\delta]$, for a small δ). The solution of this problem is discussed in the next section.

3.3.4 Quasi-convex Optimization

The optimization problem in (37) is still non-convex. It is first converted into a quasi-convex form as follows,

$$
\begin{aligned}
&\mathbf{g}(\alpha) = \text{argmax}_{\mathbf{g}\in\mathbf{R}^N}\left(\frac{2\mathbf{g}^T\mathbf{c}_{xy}+C_2}{\sigma_x^2+\mathbf{g}^T\mathbf{K}_{yy}\mathbf{g}+C_2}\right),\\
&\text{subject to : } \mathbf{g}^T\mathbf{e}=\alpha,\\
&\Leftrightarrow\\
&\min:\gamma\\
&\text{subject to : } \left[\begin{array}{c}\max:\left(\frac{2\mathbf{g}^T\mathbf{c}_{xy}+C_2}{\sigma_x^2+\mathbf{g}^T\mathbf{K}_{yy}\mathbf{g}+C_2}\right)\le\gamma\\ \text{subject to : } \mathbf{g}^T\mathbf{e}=\alpha,\end{array}\right]\\
&\Leftrightarrow\\
&\min:\gamma\\
&\text{subject to : } \left[\begin{array}{c}\min:[\gamma(\sigma_x^2+\mathbf{g}^T\mathbf{K}_{yy}\mathbf{g}+C_2)-(2\mathbf{g}^T\mathbf{c}_{xy}+C_2)]\ge 0\\ \text{subject to : } \mathbf{g}^T\mathbf{e}=\alpha\end{array}\right].
\end{aligned}
\tag{38}
$$

The first step involves the introduction of the auxiliary variable γ as an upper bound on (37). The first equivalence relation is true since minimizing γ is the same as finding the least upper bound of the function in (37). This is equal to the maximum value of the function, which exists, as seen by straightforward continuity arguments. The second equivalence relation holds since the denominator in (37) is strictly positive, allowing for the rearrangement of terms. γ them becomes a true upper bound if the problem,

$$
\left[\begin{array}{c}\max_{\mathbf{g}\in\mathbf{R}^N}:\gamma(\sigma_x^2+\mathbf{g}^T\mathbf{K}_{yy}\mathbf{g}+C_2)-(2\mathbf{g}^T\mathbf{c}_{xy}+C_2)\\ \text{subject to : } \mathbf{g}^T\mathbf{e}=\alpha\end{array}\right],
\tag{39}
$$

has a non-negative optimal value. The objective function is a linear term minus a convex quadratic and is therefore concave. The constraint is affine, and thus convex. Therefore, the overall problem is convex, and can be solved by introducing a Lagrange multiplier λ and applying the first order sufficiency conditions,

$$
\nabla_g\{\gamma(\sigma_x^2+\mathbf{g}^T\mathbf{K}_{yy}\mathbf{g}+C_2)-(2\mathbf{g}^T\mathbf{c}_{xy}+C_2)+\lambda(\mathbf{g}^T\mathbf{e}-\alpha)\}=0.
\tag{40}
$$

The solutions for $\mathbf{g}$ and λ, denoted by $\mathbf{g}(\alpha),\lambda(\alpha)$ to emphasize their dependence on α, are given by

$$
\begin{aligned}
\mathbf{g}(\alpha) &= \frac{1}{2\gamma}\mathbf{K}_{yy}^{-1}(2\mathbf{c}_{xy}-\lambda(\alpha)\mathbf{e})\\
\lambda(\alpha) &= \frac{1}{\mathbf{e}^T\mathbf{K}_{yy}^{-1}\mathbf{e}}(2\mathbf{c}_{xy}^T\mathbf{K}_{yy}^{-1}\mathbf{e}-2\gamma\alpha).
\end{aligned}
\tag{41}
$$

The optimal γ can then be computed in $O(\log(1/\varepsilon))$ iterations using a standard bisection procedure. The algorithm is summarized in Fig. 3.3.4. The tolerance specified by ε determines the tightness of the bound γ.

1. Pick an initial guess of γ (say γ_0) between 0 and 1. upper_limit = 1, lower_limit = γ_0.
2. Evaluate the optimal filter.
3. Is $\gamma(\sigma_x^2 + \mathbf{g}^T\mathbf{K}_{yy}\mathbf{g} + C_2) - (2\mathbf{g}^T\mathbf{c}_{xy} + C_2) \geq 0$?
   3a. If true, is (upper_limit - lower_limit < ε)?
      3aa. If true, we have found a γ within ε of the optimal value. Exit.
      3ab. If false, set γ_i = 0.5*(upper_limit + lower_limit), upper_limit = γ_i. Goto step 2.
   3b. If false, set γ_i = 0.5*(upper_limit + lower_limit), lower_limit = γ_i. Goto step 2.

Fig. 3 An algorithm to search for the optimal γ.

3.3.5 Search for α

It should be noted that the solution in (41) is still a function of α. The overall solution to (36) is found by searching over α.

One method is to simply initialize α to be the sum of the filter coefficients of the MSE-optimal filter, i.e., $\alpha_{init} = \mathbf{g}_{mse}^T\mathbf{e}$. Another heuristic method is to initialize α to be the sum of the filter coefficients of a structure-optimal filter. By *structure-optimal filter* is meant a filter that optimizes only the structure term in the StatSSIM index without any constraints on the mean. In other words, the goal is to find a filter $\mathbf{g}_{struct}^*$ such that

$$\begin{aligned}\mathbf{g}_{struct}^* = \text{argmax}_{\mathbf{g}\in\mathbf{R}^N}\text{Structure}(x[n],\hat{x}[n]) &= \left(\frac{2E[(x[n]-\mu_x)(\hat{x}[n]-\mu_{\hat{x}})]+C_2}{E[(x[n]-\mu_x)^2]+E[(\hat{x}[n]-\mu_{\hat{x}})^2]+C_2}\right)\\ &= \left(\frac{2\mathbf{g}^T\mathbf{c}_{xy}+C_2}{\sigma_x^2+\mathbf{g}^T\mathbf{K}_{yy}\mathbf{g}+C_2}\right).\end{aligned} \tag{42}$$

This problem is very similar to (37) and can be solved using the technique described above. The solution is

$$\mathbf{g}_{struct}^* = \frac{1}{\gamma_{struct}}(\mathbf{K}_{yy})^{-1}\mathbf{c}_{xy}, \tag{43}$$

and so the initial value of α is

$$\alpha_{init} = \mathbf{e}^T\mathbf{g}_{struct}^*. \tag{44}$$

The value of γ_{struct} is computed using the algorithm described in Section 3.3.4.

3.3.6 Application to Image Denoising and Restoration

The StatSSIM-optimal equalizer can be applied to restore images as outlined in the following steps.

- At each pixel in the distorted image, estimate the values of $\mathbf{r}_{xy}, \mathbf{c}_{xy}, \mathbf{R}_{yy}, \mathbf{K}_{yy}$ from a neighborhood of size $L \times L$. The value of L is chosen so as to compute stable correlation values.
- The minimum MSE solution is computed as $\mathbf{g}^*_{mse} = \mathbf{R}_{yy}^{-1}\mathbf{r}_{xy}$ (after removing mean from the blocks).
- The StatSSIM-optimal solution is computed as follows,
  - α is initialized using (43).
  - In a small range around α (chosen above), compute the solution in (41) and choose the one with the maximum StatSSIM index.

As before, the blurring filter and the power spectral density of the additive noise component are assumed to be known at the receiver. The procedure used to estimate the correlation and covariance values for denoising and restoration can be found in [30]. To find the estimates, the neighborhood of size $L \times L$ is unwrapped into a vector of size $L^2 \times 1$. The results of the image restoration procedure is illustrated in Figs. 4, 5, 6 and 7. The effectiveness of the StatSSIM-optimal solution is very clear from these illustrations.

3.4 SSIM-optimal Soft-thresholding

While the previous subsection explored the optimization of a linear system with respect to the SSIM index, this subsection explores SSIM-optimal soft-thresholding, a non-linear image denoising solution. A soft-thresholding operator with threshold λ is defined as

$$g(y) = \text{sgn}(y)(|y| - \lambda)_+, \tag{45}$$

where $\text{sgn}(y)$ is the signum function and $(.)_+$ is the rectifier function.

Soft-thresholding for signal denoising was first proposed by Donoho [15, 16, 14] and has been extended to image denoising most notably by Chang et al. [9]. These are risk/cost minimizing solutions where the risk or cost function is the mean squared error between the estimated source and the noise-free source. In other words, these solutions find λ that is MSE-optimal. While these solutions were proposed over two decades ago, they continue to be relevant for image denoising. The soft-thresholding problem has been solved for SSIM optimality [11] and is discussed next.

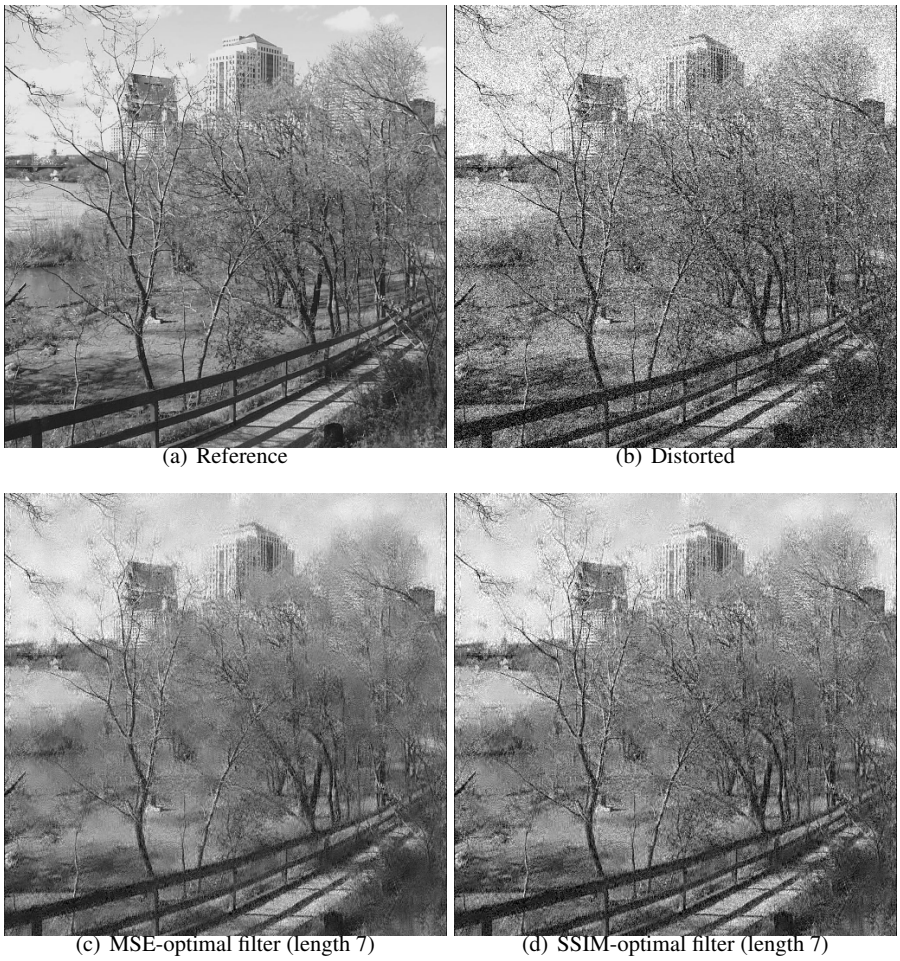

(a) Reference (b) Distorted

(c) MSE-optimal filter (length 7) (d) SSIM-optimal filter (length 7)

Fig. 4 *Denoising example 1*: Img0039.bmp from the 'City of Austin' database. 4(a) Original image. 4(b) Distorted image with $\sigma_{noise} = 35$, MSE = 1226.3729, SSIM index = 0.5511. 4(c) Image denoised with a 7-tap MSE-optimal filter, MSE = 436.6929, SSIM index = 0.6225. 4(d) Image denoised with a 7-tap SSIM-optimal filter, MSE = 528.0777, SSIM index = 0.6444.

3.4.1 SSIM index in the Wavelet Domain

The soft-thresholding problem is defined and solved in the wavelet domain while the SSIM index is defined in the space domain. To bridge this gap, the SSIM index is first expressed in the wavelet domain. This is achieved by expressing the space domain mean, variance, and cross-covariance terms in terms of wavelet coefficients. Of the several classes of wavelet transforms, only orthonormal wavelets are energy

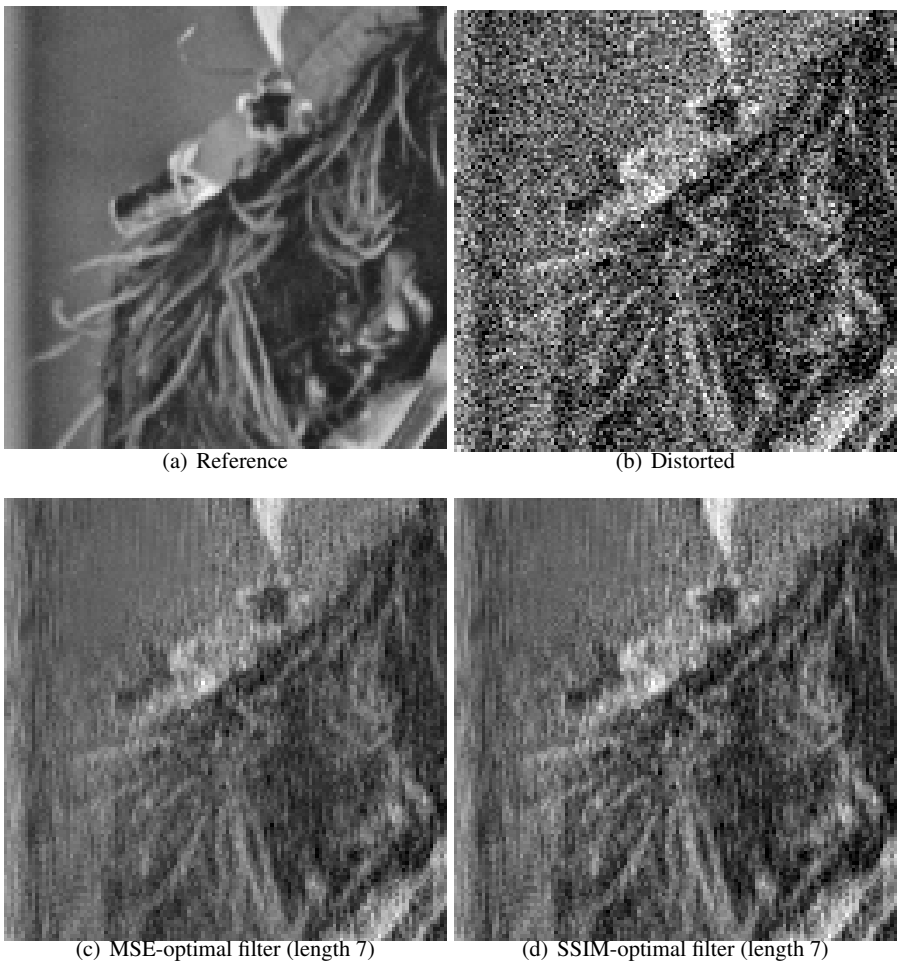

(a) Reference (b) Distorted

(c) MSE-optimal filter (length 7) (d) SSIM-optimal filter (length 7)

Fig. 5 *Denoising example 2*: 5(a) Original image. 5(b) Distorted image with $\sigma_{noise} = 40$, MSE = 1639.3132, SSIM index = 0.541485. 5(c) Image denoised with the a 7-tap MSE-optimal filter, MSE = 383.3375, SSIM index = 0.734963. 5(d) Image denoised with a 7-tap SSIM-optimal filter, MSE = 455.2577, SSIM index = 0.753917.

preserving. This property allows for the space domain variance and covariance terms to be expressed in terms of the wavelet coefficients in a straightforward manner.

The approximation subband (low-low (LL) subband) of the wavelet decomposition contains all the information required to calculate the mean of the space domain signal. A scaling factor k is applied to the mean of the LL subband to find the mean. Let $\mathbf{x}$ denote an image patch of size $N \times N$ and $\mathbf{X}$ denote the L level wavelet transform of the patch (also of size $N \times N$). Then the mean of $\mathbf{x}$ is given by

(a) Reference (b) Distorted

(c) MSE-optimal filter (d) StatSSIM-optimal filter

Fig. 6 *Restoration example 1*: Image Img0073.bmp of the 'City of Austin' database. 6(a) Original image. 6(b) Distorted image with $\sigma_{blur} = 15, \sigma_{noise} = 40$, MSE = 2264.4425, SSIM index = 0.3250. 6(c) Image restored with a 11-tap MSE-optimal filter, MSE = 955.6455, SSIM index = 0.3728. 6(d) Image restored with a 11-tap SSIM-optimal filter, MSE = 1035.0551, SSIM index = 0.4215.

$$\mu_{\mathbf{x}} = (k)^L \mu_{\mathbf{X},\mathbf{LL}}, \tag{46}$$

where k is the known scaling factor, and $\mu_{\mathbf{X},\mathbf{LL}}$ is the mean of the LL subband of $\mathbf{X}$.

The fact that orthonormal wavelets obey Parseval's Theorem is used to calculate the variance and covariance terms in the SSIM index. Let $\mathbf{x}, \mathbf{y}$ represent image patches of size $N \times N$ and $\mathbf{X}, \mathbf{Y}$ be their respective orthonormal wavelet transforms. From Parseval's Theorem, it follows that

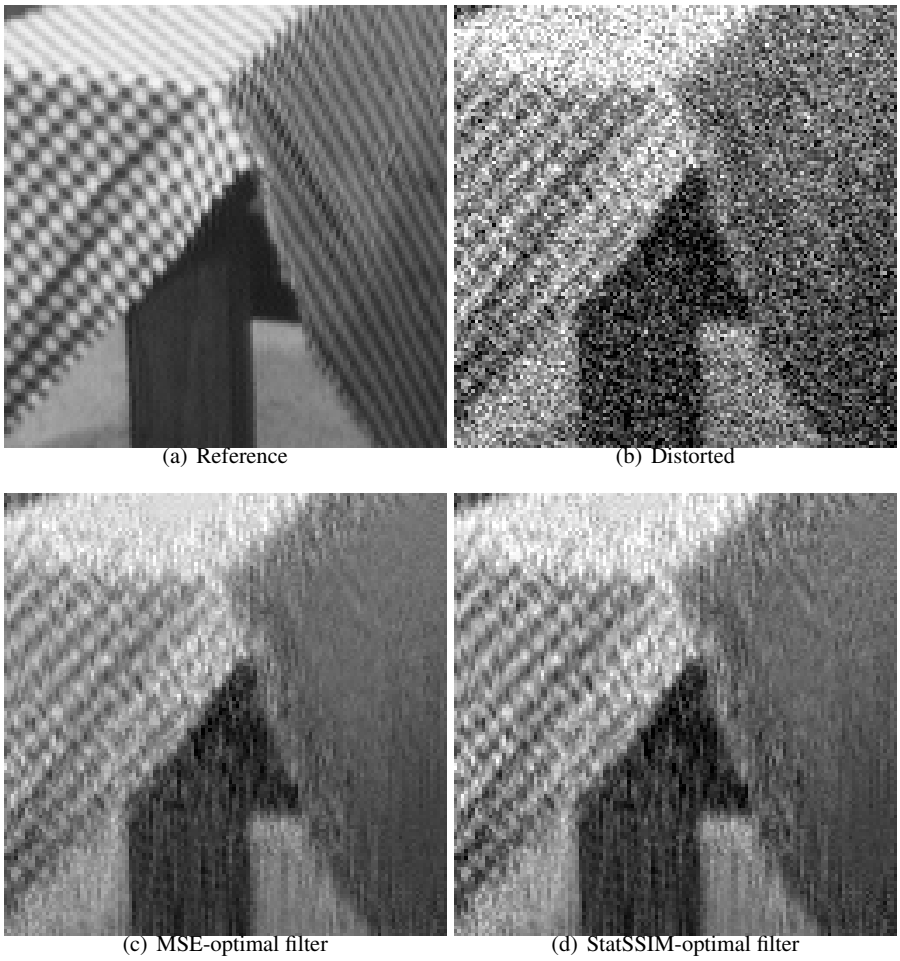

(a) Reference (b) Distorted

(c) MSE-optimal filter (d) StatSSIM-optimal filter

Fig. 7 *Restoration example 3*: A 128×128 block of Barbara image. 7(a) Original image. 7(b) Distorted image with $\sigma_{blur} = 1, \sigma_{noise} = 40$, MSE = 1781.9058, SSIM index = 0.5044. 7(c) Image restored with a 11-tap MSE-optimal filter, MSE = 520.1322, SSIM index = 0.6302. 7(d) Image restored with a 11-tap SSIM-optimal filter, MSE = 584.9232, SSIM index = 0.6568.

$$\sigma_{\mathbf{x}}^2 = \frac{1}{N^2}\sum_{i=0}^{N-1}\sum_{j=0}^{N-1} X_{i,j}^2 - ((k)^L \mu_{\mathbf{X},\mathbf{LL}})^2, \tag{47}$$

$$\sigma_{\mathbf{xy}} = \frac{1}{N^2}\sum_{i=0}^{N-1}\sum_{j=0}^{N-1} X_{i,j}Y_{i,j} - ((k)^L \mu_{\mathbf{X},\mathbf{LL}})((k)^L \mu_{\mathbf{Y},\mathbf{LL}}). \tag{48}$$

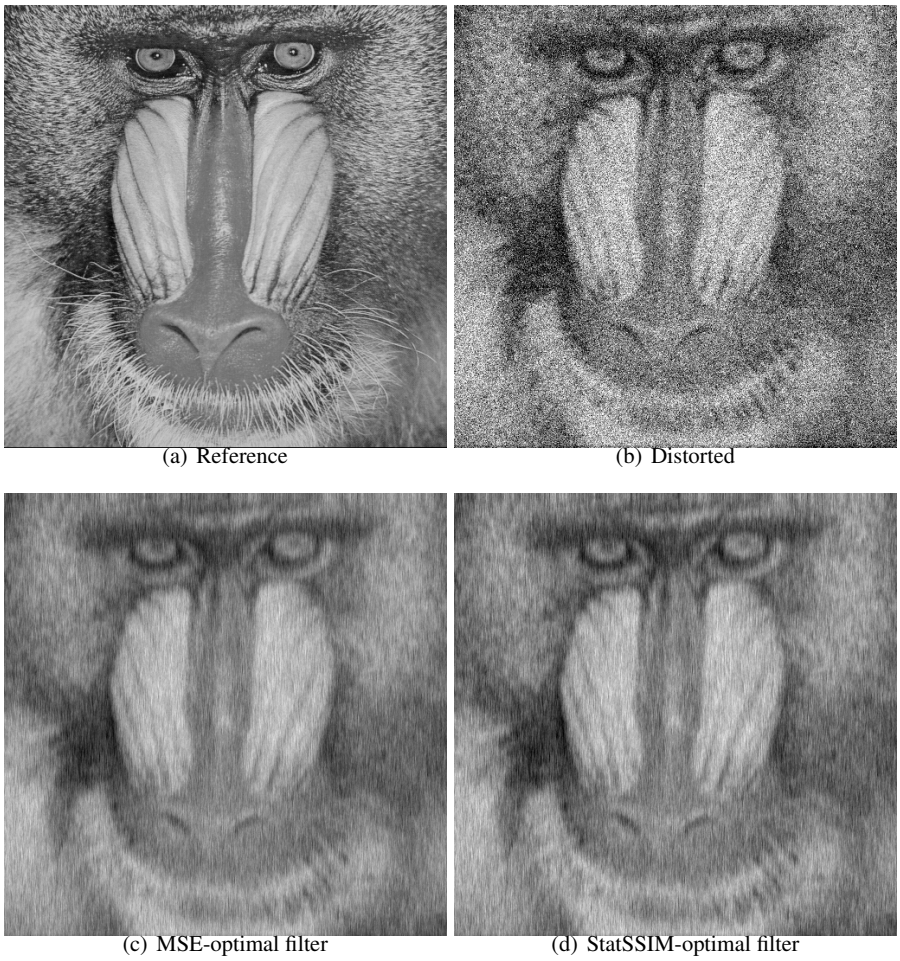

(a) Reference (b) Distorted

(c) MSE-optimal filter (d) StatSSIM-optimal filter

Fig. 8 *Restoration example 4*: A 512×512 block of Mandrill image. 8(a) Original image. 8(b) Distorted image with $\sigma_{blur} = 5, \sigma_{noise} = 50$, MSE = 3065.31, SSIM index = 0.1955. 8(c) Image restored with a 7-tap MSE-optimal filter, MSE = 863.85, SSIM index = 0.3356. 8(d) Image restored with a 7-tap SSIM-optimal filter, MSE = 908.47, SSIM index = 0.3446.

The SSIM index can now be written in terms of the wavelet coefficients as

$$SSIM(\mathbf{x},\mathbf{y}) \tag{49}$$

$$= \left(\frac{2((k)^L \mu_{\mathbf{X},\mathbf{LL}})((k)^L \mu_{\mathbf{Y},\mathbf{LL}}) + C_1}{((k)^L \mu_{\mathbf{X},\mathbf{LL}})^2 + ((k)^L \mu_{\mathbf{Y},\mathbf{LL}})^2 + C_1} \right) \left(\frac{2\frac{1}{N^2}\sum_{i=0}^{N-1}\sum_{j=0}^{N-1} X_{i,j}Y_{i,j} - ((k)^L \mu_{\mathbf{X},\mathbf{LL}})((k)^L \mu_{\mathbf{Y},\mathbf{LL}}) + C_2}{\frac{1}{N^2}\sum_{i=0}^{N-1}\sum_{j=0}^{N-1} X_{i,j}^2 + Y_{i,j}^2 - k^{2L}(\mu_{\mathbf{X},\mathbf{LL}}^2 + \mu_{\mathbf{Y},\mathbf{LL}}^2) + C_2} \right).$$

3.4.2 Problem Formulation

The soft-thresholding problem is now formulated in terms of the wavelet domain SSIM index defined in (49). Let $\mathbf{x}$ denote a pristine image patch of size $N \times N$, $\mathbf{n}$ be zero mean Gaussian noise, and $\mathbf{y} = \mathbf{x} + \mathbf{n}$ be the noisy observation of $\mathbf{x}$. Let $\mathbf{X}, \mathbf{Y}$ represent an L level orthonormal wavelet transform of $\mathbf{x}, \mathbf{y}$ respectively. Since an L level orthogonal transform consists of $3L$ subbands it leads to the design of $3L$ thresholds (one per subband). Let the threshold vector be denoted by $\Lambda = [\lambda_1, \lambda_2, \ldots, \lambda_{3L}]$ where each element corresponds to the threshold of one subband. It should be noted that the approximation band is not thresholded. Let $\hat{\mathbf{X}}$ be the soft thresholded output, and let $\hat{\mathbf{x}}$ denote the space domain version of $\hat{\mathbf{X}}$.

As with the StatSSIM-optimal equalizer design, it is assumed that the noise variance is known to the receiver. The receiver has only the observation $\mathbf{y}$ and therefore a direct evaluation of the SSIM index between $\mathbf{x}$ and $\hat{\mathbf{x}}$ is not possible. In order to estimate the SSIM index, a Gaussian source model is applied to the pristine wavelet coefficients. Since the noise is assumed to be zero mean, the mean of the pristine image patch and the thresholded estimate are identical (since the approximation subband is not thresholded). This results in the mean term of the SSIM index becoming unity. Further, since the noise is additive, the source variance can be estimated to be the difference between the variance of the observation $\sigma_{\mathbf{y}}^2$ and the noise variance $\sigma_{\mathbf{n}}^2$ as

$$\sigma_{\mathbf{x}}^2 \approx \sigma_{\mathbf{y}}^2 - \sigma_{\mathbf{n}}^2 \tag{50}$$

$$= \frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} Y_{i,j}^2 - (k^L \mu_{\mathbf{Y},LL})^2 - \sigma_{\mathbf{n}}^2. \tag{51}$$

The SSIM index is now expressed as

$$SSIM(\mathbf{x}, \hat{\mathbf{x}}) = \left(\frac{2 \frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} X_{i,j} \hat{X}_{i,j} - (k^L \mu_{\mathbf{Y},LL})^2 + C_2}{\frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} Y_{i,j}^2 + \hat{X}_{i,j}^2 - 2(k^L \mu_{\mathbf{Y},LL})^2 - \sigma_{\mathbf{n}}^2 + C_2} \right). \tag{52}$$

It should be noted that since the noise is additive, $\frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} X_{i,j} \hat{X}_{i,j} \approx (\sigma_Y^2 - \sigma_n^2 + \mu_y^2)$.

The SSIM-optimal soft-thresholding problem is formulated as:

$$\Lambda^* = \text{argmax}_{\Lambda \in \mathbf{R}_+^{3L}} SSIM(\mathbf{x}, \hat{\mathbf{x}}) \tag{53}$$

3.4.3 Solution

The objective function is nonlinear in Λ and maps a $3L$-dimensional vector to a one-dimensional scalar. The optimization is constrained by the requirement for Λ to be non-negative. The quasi-Newton optimization method provides a good tradeoff between complexity and performance in finding local optima. The Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm [3] is one such method that was employed here to find local optima. As a consequence, this solution is locally optimal and no guarantees on global optimality can be made. The application of this solution to the image data is presented in the following steps.

- The noisy image is divided into non-overlapping blocks of size 32×32
- Apply a L level orthonormal wavelet transform to each block
- For each wavelet transformed block, compute its subband statistics
- Find a locally optimal Λ^* using the BFGS algorithm. The search can be initialized to the MSE-optimal soft thresholding solution from Chang et al. [9].
- Using Λ^* to soft-threshold the corresponding wavelet subband coefficients
- The denoised coefficients are then transformed back to the pixel domain by applying the inverse wavelet transform on a block by block basis

To perform a qualitative and quantitative comparison, the locally SSIM-optimal solution is compared with the MSE-optimal solution by Chang et al. [9]. The denoising results are presented in Fig. 9. It can be observed that the SSIM-based solution retains more image detail and has a better perceptual quality compared to the MSE-optimal solution. This method provides conclusive evidence that optimization for perceptual quality is a worthwhile endeavor indeed.

4 Local SSIM-optimal Approximation

We now study whether a perceptual criterion such as SSIM can be used for best-basis approximation. This section follows the treatment in [5].

The Minimal Mean Square Error approximation problem is

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{z} \in A} \|\mathbf{x} - \mathbf{z}\|_2^2, \tag{54}$$

where $\mathbf{x}$ is the image patch to be approximated, $\hat{\mathbf{x}}$ is the best approximation and A is a constraint that represents some prior knowledge on $\mathbf{x}$.

A very successful and popular image prior is based on sparse representation. Sparsity is the image model that assumes that any image patch can be represented by a linear combination of a few atomic elements. If the matrix Ψ is a dictionary of basis elements that is multiplied by coefficients $\mathbf{c}$, then a sparsity constraint is $\|\mathbf{c}\|_0 \leq K$, which means that there are K or fewer non-zero elements of $\mathbf{c}$. We enumerate the columns of Ψ as $(\psi_1, \psi_2, \ldots, \psi_P)$. The set A is then the span of Ψ:

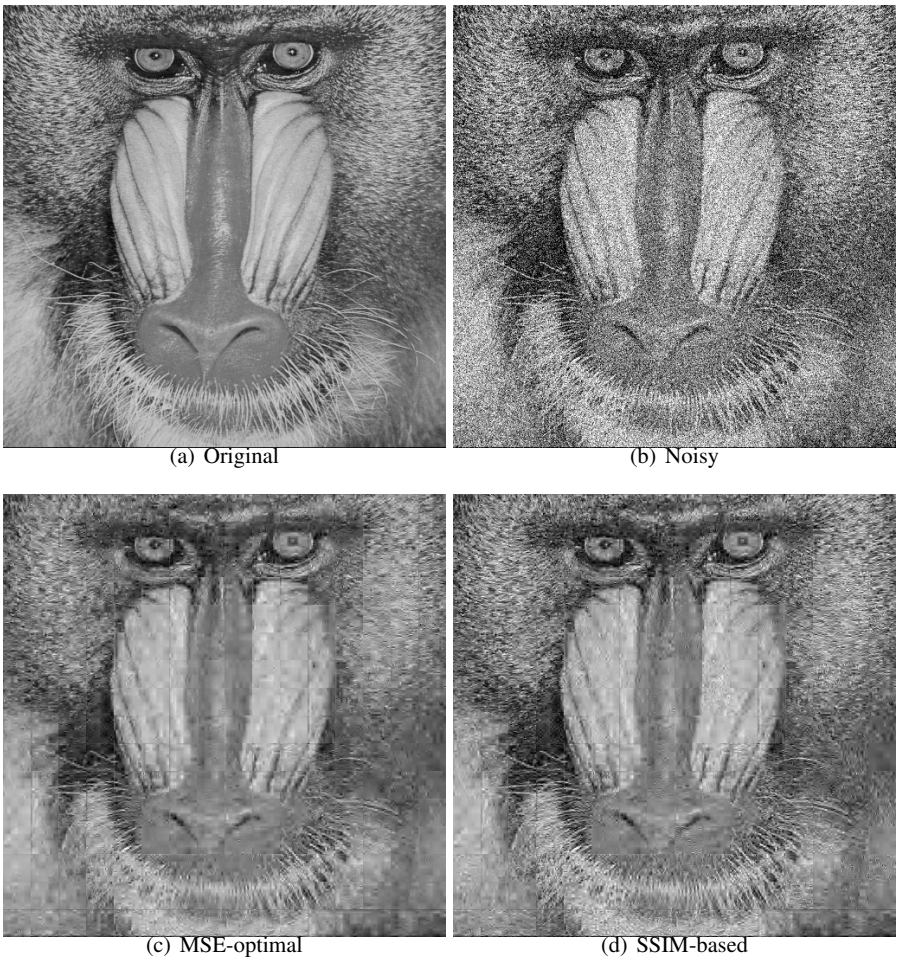

(a) Original (b) Noisy

(c) MSE-optimal (d) SSIM-based

Fig. 9 9(a) Undistorted Mandrill image. 9(b) Noisy image with $\sigma_n = 50$, MSE = 2492.49, SSIM index = 0.2766. 9(c) MSE-optimal soft thresholding, MSE = 509.16, SSIM index = 0.4835. 9(d) SSIM-based soft thresholding, MSE = 577.54, SSIM index = 0.4954.

$$A = \text{span}\{\psi_1, \psi_2, \cdots, \psi_P\}. \tag{55}$$

The optimization of the mean square error can be interpreted in two different ways: either 1) a perfect approximation is assumed, except for a stochastic additive white Gaussian noise part, or 2) an imperfect approximation is assumed, but with the mean squared error as a model of dissimilarity. As it is empirically observed that images are not exactly sparse, but more accurately compressible, that is almost sparse, the first assumption cannot be valid. Moreover, as the mean squared error is a poor model of perception, the second assumption should also be dismissed.

The best approximation problem should instead be posed in term of a perceptual measure. We solve this problem below for the SSIM index case. But before embarking into SSIM-based approximation, it will be useful to recall L^2-based approximation results. As we will see later on, the two problems share striking characteristics.

4.1 L^2-based Approximation

The solution of the L^2-based sparse approximation problem can be divided into three cases: orthogonal basis, linear redundant basis and non-linear approximation.

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

$$\mathbf{x} = \sum_{k=1}^{N} a_k \psi_k, \quad a_k = \psi_k^T \mathbf{x}, \;\; 1 \le k \le N. \tag{56}$$

The expansions of the approximation $\hat{\mathbf{x}}$ will be denoted as

$$\hat{\mathbf{x}} = \sum_{k=1}^{N} c_k \psi_k, \tag{57}$$

where $(c_1, c_2, \ldots, c_N)$ are unknown coefficients.

4.1.1 Orthogonal Basis

In the orthogonal case, the approximation spaces A in (54) will be the span of subsets of the set of basis functions $\{\psi_k\}_{k=1}^N$. At this point, we do not exactly specify which other ψ_k basis functions will be used but consider all possible subsets of $M < N$ basis functions:

$$A = \text{span}\{\psi_{\gamma(1)}, \psi_{\gamma(2)}, \cdots, \psi_{\gamma(M)}\}, \tag{58}$$

where $\gamma(i) \in \{1, 2, \cdots, N\}$ and $c_{\gamma(M+1)} = \cdots = c_{\gamma(N)} = 0$.

The well-known L^2-based optimal approximation is summarized in the following theorem:

Theorem 4.1 *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 the M Fourier coefficients $a_k = \langle \mathbf{x}, \psi_k \rangle$ of greatest magnitude, i.e.*

$$c_{\gamma(k)} = \begin{cases} a_{\gamma(k)} = \psi_{\gamma(k)}^T \mathbf{x}, & 1 \le k \le M, \\ 0, & M+1 \le k \le N. \end{cases} \tag{59}$$

where $|a_{\gamma(1)}| \ge |a_{\gamma(2)}| \ge \ldots \ge |a_{\gamma(M)}| \ge |a_l|$ with $l \in \{1, 2, \cdots, N\} \setminus \{\gamma(1), \cdots, \gamma(M)\}$.

4.1.2 Linear Redundant Basis

A redundant or over-complete basis of $\mathbf{R}^N$ consists of $P \geq N$ column vectors $\Psi = \{\psi_k\}_{k=1}^P$ such that N of them are linearly independent. Given $\mathbf{x}$ in $\mathbf{R}^N$, we search for an approximation $\hat{\mathbf{x}}$ of $\mathbf{x}$ with the help of the first $M < N$ linearly independent vectors of a redundant basis:

$$\hat{\mathbf{x}} = \sum_{k=1}^{M} c_k \psi_k. \tag{60}$$

We seek the coefficients $\mathbf{c} = [c_1, c_2, \ldots, c_M]$ that will minimize the L^2-error:

$$\|\mathbf{x} - \hat{\mathbf{x}}\|_2^2 = (\mathbf{x} - \sum_{j=1}^{M} c_j \psi_j)^T (\mathbf{x} - \sum_{k=1}^{M} c_k \psi_k) \tag{61}$$

The solution is the well-known normal equation,

$$\sum_{j=1}^{M} c_j \psi_j^T \psi_k = \psi_k^T \mathbf{x}, \quad 1 \leq k \leq M. \tag{62}$$

Thus, to find the optimal coefficients $\mathbf{c}$, we need to solve a $M \times M$ linear system of equations $\Phi\mathbf{c} = \mathbf{a}$ with $\phi_{j,k} := \psi_j^T \psi_k$ and $\mathbf{a} = [a_1, a_2, \ldots, a_M]$ where $a_k = \psi_k^T \mathbf{x}$. In practice, $\mathbf{c}$ is found by multiplying the pseudo-inverse Ψ^+ of the dictionary matrix Ψ with $\mathbf{x}$:

$$\mathbf{c} = (\Psi^T \Psi)^{-1} \Psi^T \mathbf{x} \tag{63}$$

$$= \Psi^+ \mathbf{x}. \tag{64}$$

This pseudo-inverse can be computed from the singular value decomposition $U\Sigma V^T$ of Ψ:

$$\Psi^+ = V \Sigma^+ U^T, \tag{65}$$

where Σ^+ is the diagonal matrix whose positive elements are the reciprocal of the non-zero elements of Σ.

4.1.3 Non-Linear Approximation

The problem can be written as

$$\arg\min \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2 \quad \text{subject to } \|\mathbf{c}\|_0 \leq K, \tag{66}$$

where $\mathbf{x}$ is the signal to be approximated,

$$\hat{\mathbf{x}} := \Psi\mathbf{c} = \sum_{k=1}^{P} c_k \psi_k \tag{67}$$

is the approximation and

$$\|\mathbf{c}\|_0 := \sum_{k=1}^{P} c_k^0 = \#\{k : |c_k| > 0\} \tag{68}$$

is the 0-pseudonorm. We follow the convention that $0^0 = 0$. Thus for a non-linear approximation, we choose the M best vectors from the $P \geq N$ vectors $\Psi = \{\psi_k\}_{k=1}^{P}$. We solve a linear system in a manner similar to the linear case (see (62)) in order to determine the coefficients $\mathbf{c}$. There are $\frac{P!}{(P-M)!M!}$ possibilities, which grow exponentially with P. In fact, it has been shown that finding the sparse approximation that minimizes $\|\mathbf{x} - \hat{\mathbf{x}}\|_2^2$ is an NP-hard problem [13].

Two approaches have been adopted to avoid the NP-hard problem: Matching Pursuit (MP) and Basis Pursuit.

Matching Pursuit

The MP algorithm from Mallat and Zhang [20] greedily adds vectors one at a time until a M- vector approximation is found. For the first vector, we minimize

$$\|\mathbf{x} - c_{\gamma(1)}\psi_{\gamma(1)}\|_2^2 = \|\mathbf{x}\|_2^2 - 2c_{\gamma(1)}a_{\gamma(1)} + a_{\gamma(1)}^2. \tag{69}$$

By taking partial derivatives and setting them to zero, we find that the solution is exactly the same as the orthogonal case. We choose the index that maximizes $|a_m| = |\psi_m^T\mathbf{x}|$.

For the K-th vector, the L^2-error is

$$\begin{aligned}\|\mathbf{x} - \sum_{k=1}^{K} c_{\gamma(k)}\psi_{\gamma(k)}\|_2^2 = \|\mathbf{x} - \sum_{k=1}^{K-1} c_{\gamma(k)}\psi_{\gamma(k)}\|_2^2 + c_{\gamma(K)}^2\psi_{\gamma(K)}^T\psi_{\gamma(K)} \\ -2c_{\gamma(K)}\psi_{\gamma(K)}^T(\mathbf{x} - \sum_{k=1}^{K-1} c_{\gamma(k)}\psi_{\gamma(k)}).\end{aligned} \tag{70}$$

The error will be minimized when $|\psi_m^T(\mathbf{x} - \sum_{k=1}^{K-1} c_{\gamma(k)}\psi_{\gamma(k)})|$ is maximized. Note that in the orthogonal basis case, the matching pursuit algorithm coincides with the optimal algorithm which chooses the M largest basis coefficients.

Orthogonal Matching Pursuit

The Orthogonal Matching Pursuit (OMP) [26] combines the MP with a Gram-Schmidt procedure in order to obtain an orthogonal basis.

Given a linearly independent basis $\{\psi_1, \psi_2, \ldots, \psi_N\}$ of $\mathbf{R}^N$, the Gram-Schmidt procedure successively projects the basis to an orthogonal subspace

$$G_1 = \psi_1 \tag{71}$$

$$G_2 = \psi_2 - \frac{G_1^T \psi_2}{G_1^T G_1} G_1 \tag{72}$$

$$\vdots$$

$$G_N = \psi_N - \sum_{j=1}^{N-1} \frac{G_j^T \psi_N}{G_j^T G_j} G_j. \tag{73}$$

The basis is then normalized with $g_k = G_k / \|G_k\|_2$.

The OMP algorithm thus alternates between finding the best matching vector and the orthonormalization process. The trade-off is a convergence with a finite number of iterations against an extra computational cost for the orthonormalization.

For the numerical implementation, the naive approach of computing an orthonormal basis is well known to be unstable. A more reliable way to perform the OMP is with the following algorithm [33]:

1. Initialize: $I = \varnothing, \mathbf{r} := \mathbf{x}$ and $\mathbf{a} := \mathbf{0}$.
2. While $\|\mathbf{r}\|^2 > T$, do
3. $k^* := \arg\max_k |\psi_k^T \mathbf{r}|$;
4. Add k^* to the set of indices I;
5. $\mathbf{a}_I := \Psi_I^+ \mathbf{x}$;
6. $\mathbf{r} := \mathbf{x} - \Psi_I \mathbf{a}_I$;
7. end while.

Here, $\mathbf{a}_I$ and Ψ_I represent the restriction of, respectively, $\mathbf{a}$ and Ψ to the elements or columns of indices I. Note that since the pseudo-inverse of Ψ has to be computed for incrementally larger matrices, there are ways to make computations more efficient with the help of Cholesky factorization (see [33]).

It is not immediately clear whether this algorithm really performs the OMP. To see this, note that since $\mathbf{g} = [g_1, g_2, \ldots, g_K]$ is an orthonormal basis, the change of basis,

$$\sum_{j=1}^{K} a_j \psi_j = \sum_{j=1}^{K} g_j^T \mathbf{x} g_j, \tag{74}$$

is performed via

$$\mathbf{a} = (\Psi^T \Psi)^{-1} \Psi^T \mathbf{g} \mathbf{g}^T \mathbf{x} \tag{75}$$

$$= \Psi^+ \mathbf{x}. \tag{76}$$

Thus, the residual $\mathbf{r}$ can be computed from the original basis Ψ and the orthonormalization is hidden in the computation of the coefficients $\mathbf{a}$.

Basis Pursuit

In Basis Pursuit, the non-linear approximation problem is replaced by an L^1-regularization problem:

$$\arg\min \|\mathbf{x}-\hat{\mathbf{x}}\|_2^2 \quad \text{subject to } \|\mathbf{c}\|_1 \leq T, \tag{77}$$

for some constant T. The choice of the L^1-norm makes the problem convex. Details on techniques to solve this problem are given in [12].

4.2 SSIM-based Approximation

The problem now can be stated as follows:

Given $\mathbf{x}$ and Ψ, find the $\mathbf{z} = \Psi\mathbf{c}$ that maximizes $S := \mathrm{SSIM}(\mathbf{x}, \mathbf{z})$.

This problem was first solved for the orthogonal case in [6] before being generalized to the redundant case in [32] and [31].

4.2.1 Linear Approximation

In a linear approximation, the choice of dictionary vectors $\Psi = (\psi_1, \psi_2, \ldots, \psi_M)$ is already fixed and we need only to find the coefficients $\mathbf{c} = (c_1, c_2, \ldots, c_M)$ that maximize the SSIM. To do that, we search for the stationary points of the partial derivatives of SSIM with respect to c_k. First, we write the mean, the variance and the covariance of $\mathbf{z}$ in terms of $\mathbf{c}$:

$$\mu_z = \sum_{k=1}^{M} c_k \bar{\psi}_k, \tag{78}$$

$$(N-1)\sigma_z^2 = \sum_{j=1}^{M}\sum_{k=1}^{M} c_j c_k \psi_j^T \psi_k - N\mu_z^2,$$

$$(N-1)\sigma_{x,z} = \sum_{k=1}^{M} c_k \psi_k^T \mathbf{x} - N\mu_x\mu_z. \tag{79}$$

Next, we find the partial derivatives:

$$\frac{\partial \mu_z}{\partial c_k} = \bar{\psi}_k; \tag{80}$$

$$(N-1)\frac{\partial \sigma_z^2}{\partial c_k} = 2\sum_{j=1}^{M} c_j \psi_j^T \psi_k - 2N\mu_x\bar{\psi}_k; \tag{81}$$

$$(N-1)\frac{\partial \sigma_{x,z}}{\partial c_k} = \psi_k^T \mathbf{x} - N\mu_x\bar{\psi}_k. \tag{82}$$

The logarithm of SSIM can be written as

$$\begin{aligned}\log S = {} & \log(2\mu_x\mu_z + C_1) - \log(\mu_x^2 + \mu_z^2 + C_1) \\ & + \log(2\sigma_{x,z} + C_2) - \log(\sigma_x^2 + \sigma_z^2 + C_2).\end{aligned} \tag{83}$$

So for all $1 \le k \le M$,

$$\begin{aligned}\frac{1}{S}\frac{\partial S}{\partial c_k} = {} & \frac{2\mu_x\bar{\psi}_k}{2\mu_x\mu_z + C_1} - \frac{2\mu_z\bar{\psi}_k}{\mu_x^2 + \mu_z^2 + C_1} \\ & + \frac{2\psi_k^T\mathbf{x} - 2N\bar{x}\bar{\psi}_k}{(N-1)(2\sigma_{x,z} + C_2)} - \frac{2\sum_{j=1}^M c_j\psi_j^T\psi_k - 2N\mu_z\bar{\psi}_k}{(N-1)(\sigma_x^2 + s_z^2 + C_2)}.\end{aligned} \tag{84}$$

Solution for Oscillatory Basis

In the particular case where the basis is comprised of normalized oscillatory functions, we have

$$\bar{\psi}_k = 0 \text{ and } \|\psi_k\| = 1 \quad \text{for } 1 \le k \le M. \tag{85}$$

This leads to

$$\mu_z = 0, \tag{86}$$

$$(N-1)\sigma_z^2 = \sum_{j=1}^M \sum_{k=1}^M c_j c_k \psi_j^T \psi_k \tag{87}$$

$$(N-1)\sigma_{x,z} = \sum_{k=1}^M c_k \psi_k^T \mathbf{x}. \tag{88}$$

Therefore the partial derivative in (84) becomes

$$\frac{\partial S}{\partial c_k} = S\left[\frac{2\psi_k^T\mathbf{x}}{(N-1)(2\sigma_{x,z} + C_2)} - \frac{2\sum_{j=1}^M c_j\psi_j^T\psi_k}{(N-1)(\sigma_x^2 + \sigma_z^2 + C_2)}\right]. \tag{89}$$

We now search for stationary points:

$$\frac{\partial S}{\partial c_k} = 0 \Rightarrow \frac{\psi_k^T\mathbf{x}}{\sum_{j=1}^K c_j\psi_j^T\psi_k} = \frac{2\sigma_{x,z} + C_2}{\sigma_x^2 + \sigma_z^2 + C_2} =: \frac{1}{\alpha}, \quad \text{for } 1 \le k \le M. \tag{90}$$

We can rewrite this equation as

$$\sum_{j=1}^M c_j\psi_j^T\psi_k = \alpha\psi_k^T\mathbf{x}, \quad \text{for } 1 \le k \le M. \tag{91}$$

This equation is very similar to (62) for the optimal coefficients for the L^2-based approximation. In fact, since the equations (62) and (91) are identical up to a scaling factor and since the solution of the linear system is unique, we have

$$c_k = \alpha a_k. \tag{92}$$

Now, we seek to find an expression for α. Starting from the right hand side of (90), we replace σ_z^2 and $\sigma_{x,z}$ by their basis expansion (86) and then employ (92) for the c_k to obtain

$$\alpha = \frac{\alpha^2 A + B}{\alpha C + D}, \tag{93}$$

where

$$A = \frac{1}{N-1} \sum_{j=1}^{M} \sum_{k=1}^{M} a_j a_k \psi_j^T \psi_k, \tag{94}$$

$$B = \sigma_x^2 + C_2, \tag{95}$$

$$C = \frac{2}{N-1} \sum_{k=1}^{M} a_k^2, \tag{96}$$

$$D = C_2. \tag{97}$$

Equation (93) is a quadratic equation in α with solutions,

$$\alpha = \frac{-D \pm \sqrt{D^2 + 4(C-A)B}}{2(C-A)} \tag{98}$$

$$= \frac{-C_2 \pm \sqrt{C_2^2 + 2C(\sigma_x^2 + C_2)}}{C}. \tag{99}$$

Note that $C - A = C/2 = A$, since the a_k's are found by solving the linear system (62).

Flat Approximation Case

We now consider the case in which a flat basis function $\psi_0 \equiv 1$ is added to the oscillatory basis. In this case,

$$\mu_z = c_0. \tag{100}$$

The coefficient c_0 is the stationary point of (84):

$$\frac{\partial S}{\partial c_0} = S\left[\frac{2\mu_x\bar{\psi}_0}{2\mu_x c_0+C_1} - \frac{2c_0\bar{\psi}_0}{\mu_x^2+c_0^2+C_1} + \frac{2\psi_0^T\mathbf{x}-2N\mu_x\bar{\psi}_0}{(N-1)(2\sigma_{x,z}+C_2)} - \frac{2\sum_{j=0}^{M}c_j\psi_j^T\psi_0-2Nc_0\bar{\psi}_0}{(N-1)(\sigma_x^2+\sigma_z^2+C_2)}\right]$$
$$= S\left[\frac{2\mu_x}{2\mu_x c_0+C_1} - \frac{2c_0}{\mu_x^2+c_0^2+C_1}\right]. \tag{101}$$

Solving for the stationary point leads to the following quadratic equation in c_0:

$$c_0^2\bar{x}+C_1c_0-\mu_x(\mu_x^2+C_1)=0. \tag{102}$$

Its solution is

$$c_0 = \frac{-C_1 \pm \sqrt{C_1^2+4\mu_x^2(\mu_x^2+C_1)}}{2\mu_x}. \tag{103}$$

We choose the positive branch to maximize the SSIM index, which is simply $c_0 = \mu_x$, as expected. The other coefficients are found as in the oscillatory basis case.

Orthogonal Basis

In the orthogonal case, the constants in the equation for α (99) simplify to

$$\alpha = \frac{-C_2 \pm \sqrt{C_2^2+(\frac{4}{N-1}\sum_{k=1}^{M}(\psi_k^T\mathbf{x})^2)(\sigma_x^2+C_2)}}{\frac{2}{N-1}\sum_{k=1}^{M}(\psi_k^T\mathbf{x})^2}. \tag{104}$$

The SSIM index is maximized with the positive branch. If $C_2 = 0$, then

$$\alpha = \frac{s_{\mathbf{x}}}{\sqrt{\frac{1}{N-1}\sum_{k=1}^{M}(\psi_k^T\mathbf{x})^2}}, \tag{105}$$
$$= \left(\frac{\sum_{k=1}^{N}a_k^2}{\sum_{k=1}^{M}a_k^2}\right)^{1/2}, \tag{106}$$

where the second equality follows from Parseval's Theorem. Thus the coefficients are adjusted in order to preserve the variance of the original signal.

4.3 Non-Linear Approximation

Similar to the L^2-case, the problem is to find, given a dictionary $\Psi \in \mathbf{R}^{P\times N}$ and a signal $\mathbf{x} \in \mathbf{R}^N$, the coefficients $\mathbf{c} \in \mathbf{R}^P$ with $\|\mathbf{c}\|_0 = M < N$ such that

$$\mathrm{SSIM}(\mathbf{x}, \Psi\mathbf{c}) \tag{107}$$

is maximized.

SSIM-based matching pursuit

From Section 4.2.1, we know how to find the best coefficients given a set of vectors. It remains to determine which vectors to choose. We assume an oscillatory dictionary which contains a flat element $\psi_0 = \mathbf{e}$. This flat element is always included in the approximation so that, by (103), $c_0 = \mu_x$ and $S_1(\mathbf{x}, \mathbf{z}) = 1$. It remains to optimize S_2.

First, we want to find ψ_{γ_0} and c_{γ_0} that will maximize $S_2(\mathbf{x}, c_{\gamma_0}\psi_{\gamma_0})$. The second component of the simplified SSIM index is written as

$$S_2(\mathbf{x}, c_{\gamma_0}\psi_{\gamma_0}) = \frac{2c_{\gamma_0}\psi_{\gamma_0}^T(\mathbf{x} - \mu_x\mathbf{e}) + C_2(N-1)}{\|\mathbf{x} - \mu_x\mathbf{e}\|^2 + c_{\gamma_0}^2 + C_2(N-1)}. \tag{108}$$

For any fixed c_k, the SSIM will be maximized when $|\psi_k^T(\mathbf{x} - \mu_x\mathbf{e})| = |\psi_k^T\mathbf{x}|$ is maximized. We thus choose

$$\gamma_0 = \arg\max_{1\le k\le P} |\psi_k^T\mathbf{x}| \tag{109}$$

and

$$c_{\gamma_0} = \alpha\psi_{\gamma_0}^T\mathbf{x}. \tag{110}$$

In general, we want to find ψ_{γ_K} and c_{γ_K} that will maximize

$$S_2(\mathbf{x}, \sum_{k=0}^{K-1} c_{\gamma_k}\psi_{\gamma_k} + c_{\gamma_K}\psi_{\gamma_K}). \tag{111}$$

For every choice of ψ_{γ_K}, we would need to find $\{a_{\gamma_k}\}_{0\le k\le K}$, i.e. we have to solve a $K \times K$ linear system of equations and compute the SSIM with $c_{\gamma_k} = \alpha a_{\gamma_k}$, then pick the basis ψ_{γ_K} that yields the maximum value. In practice this procedure is intractable given that a potentially large linear system has to be solved for every possible basis of the dictionary and at every iteration of the greedy algorithm.

SSIM-based Orthogonal Matching Pursuit

According to the MP algorithm, the choice of the first basis that maximizes the SSIM index is the same as that of the optimal L^2-basis. Indeed, (108) is maximized when $|\psi_k^T\mathbf{x}|$ is maximized.

For the choice of the K-th basis, we seek to maximize

$$S_2(\mathbf{x}, \sum_{j=1}^{K-1} c_{\gamma_j}\psi_{\gamma_j} + c_{\gamma_K}\psi_{\gamma_K}) = S(\mathbf{x}, \sum_{j=1}^{M-1} c_{\gamma_j} g_{\gamma_j} + c_{\gamma_M} g_{\gamma_M}) \tag{112}$$

$$= \frac{2\sum_{j=1}^{K-1}(g_{\gamma_j}^T\mathbf{x})c_{\gamma_j} + (g_{\gamma_K}^T\mathbf{x})c_{\gamma_K} + (N-1)C_2}{\|\mathbf{x}-\mu_x\mathbf{e}\|^2 + \sum_{j=1}^{K} c_{\gamma_j}^2 + (N-1)C_2}. \tag{113}$$

The choice of basis that will maximize the SSIM index is

$$\gamma_K = \arg\max_{1\le k\le P} |g_k^T\mathbf{x}|. \tag{114}$$

Note that

$$g_k^T\mathbf{x} = (\psi_k - \sum_{j=1}^{K-1}(\psi_k^T g_j)g_j)^T\mathbf{x} \tag{115}$$

$$= \psi_k^T(\mathbf{x} - \sum_{j=1}^{K-1}(\psi_k^T g_j)g_j) \tag{116}$$

$$= \psi_k^T\mathbf{r}. \tag{117}$$

Thus, the optimal basis for the SSIM-based and the L^2-based algorithms are exactly the same. Indeed, the SSIM-based coefficients will be simply a scaling of the L^2-based coefficients. The difference will be in the stopping criterion: the SSIM-OMP stopping criterion will depend on the SSIM index instead of the L^2-error.

4.4 Variational SSIM

Otero et al. [24, 25, 23] have explored the direction of SSIM-optimal algorithm design directly in the pixel domain. They exploit the quasi-convexity properties for zero-mean signals (see Section 2). Based on these observations, several optimization problems are formulated and solved using a simple bisection method. Because of space limitations, we omit a detailed discussion of these methods in this chapter and simply refer readers to [24, 25, 23].

5 Image-wide Variational SSIM Optimization

Block-based schemes find a perceptually optimal solution locally, but blockiness artifacts might appear when combining these local solutions to form an image. Several *ad hoc* methods such as taking the central pixel or averaging overlapping blocks are possible, but they might not be perceptually optimal.

We propose a solution based on convex or quasi-convex optimization. Let R_i be the block extraction operator, which from an image $\mathbf{x}$ gets the i-th block. Let $\mathbf{z}_i$

be the local perceptually optimal estimator corresponding to the i-th image block. Finally, let d be a quasi-convex dissimilarity measure and let A be a convex set. Then the solution of the following optimization problem is the desired image-wide perceptually optimal estimator:

$$\hat{\mathbf{x}} = \min_{\mathbf{x}\in A} \max_{i} d(R_i(\mathbf{x}), \mathbf{z}_i). \tag{118}$$

Since R_i is a linear function and d is quasi-convex, $d(R_i(\mathbf{x}), \mathbf{z}_i)$ is also quasi-convex. As summarized in Section 2, the maximum of quasi-convex is also quasi-convex. Since A is a convex set, the optimization problem can thus be solved using the bisection method.

In the special case of image denoising, A takes the form of $\|\mathbf{x}-\mathbf{y}\|_2^2 \leq \alpha$, where α is a parameter controlling how close to the noisy image $\mathbf{y}$ the restored image $\mathbf{x}$ will be. Alternatively, the estimator could be written in a variational form as

$$\hat{\mathbf{x}} = \max_{i} d(R_i(\mathbf{x}), \mathbf{z}_i) + \beta \|\mathbf{x}-\mathbf{y}\|_2^2, \tag{119}$$

where β is a parameter controlling the balance between the fit to the local perceptual estimator and the global information from the noisy input.

These optimization problems can be compared to the one defined in [31]:

$$\arg\max_{\mathbf{x}} \{SSIM(\mathbf{w}, \mathbf{x}) + \lambda SSIM(\mathbf{x}, \mathbf{y})\}. \tag{120}$$

There are two major differences: 1) the mean SSIM score is optimized instead of the minimal SSIM score, 2) the Structural Similarity between the noisy data and the global solution is taken instead of the Mean Squared Error. We argue that the new formulation (119) is not only more convenient mathematically, but also conceptually more preferable.

Indeed, it is mathematically convenient since we can prove the existence and uniqueness of a solution (assuming that d is not-flat) and we have a (bisection) method to compute the optimal solution that will always converge to the global minimum. It is also justifiable conceptually, since the maximum dissimilarity corresponds to the most salient feature, the one that will be the most perceptually annoying. (However, if the image quality is non-uniform the maximum might put too much weight on a single location.) Moreover, taking the mean-squared error between the noisy image and the restored image is justified by the model of distortion (additive noise), not the model of perception.

6 Conclusions

In this chapter, we have discussed the problem of optimizing the perceptual quality of images by employing perceptual quality measures in the optimization problem as opposed to traditional distance measures, e.g. MSE/RMSE. The SSIM index was

considered to be the cost function in this discussion. The mathematical properties of the SSIM index were first presented, most notably its quasi-convexity. Subsequently, the classical problem of image restoration was reformulated with the statistical version of SSIM index as the cost function. The solution to this problem demonstrated the gains to be had by explicitly optimizing for perceptual quality metrics. Corroborative evidence to this claim was shown in the form of a soft-thresholding solution, again optimized with respect to the SSIM index. Subsequently, a methodology for constructing SSIM-optimal basis functions was discussed and variations to popular pursuit algorithms such as matching pursuit and basis pursuit were presented. Through these discussions, a set of SSIM-optimal solutions that address popular image processing problems ranging from denoising and restoration to dictionary construction were presented. These solutions provide a platform to address a larger set of problems that could be reduced to one of these forms. Furthermore, these solutions demonstrate that optimization of perceptual quality is indeed the way forward in building next generation multimedia systems.

References

1. H. C. Andrews and B. R. Hunt, *Digital image restoration*, Prentice-Hall Signal Processing Series, Englewood Cliffs, NJ, 1977.
2. D. P. Bertsekas, *Dynamic programming and optimal control*, Athena Scientific, Belmont, MA, 1995.
3. ———, *Nonlinear programming*, Athena Scientific, Belmont, MA, 1999.
4. S. Boyd and L. Vandenberghe, *Convex optimization*, Cambridge University Press, New York, 2004.
5. D. Brunet, *A study of the structural similarity image quality measure with applications to image processing*, Ph.D. thesis, University of Waterloo, 2012.
6. D. Brunet, E. R. Vrscay, and Z. Wang, *Structural similarity-based approximation of signals and images using orthogonal bases*, Proc. Int. Conf. on Image Analysis and Recognition (M. Kamel and A. Campilho, eds.), LNCS, vol. 6111, Springer, Heidelberg, 2010, pp. 11–22.
7. ———, *On the mathematical properties of the structural similarity index*, IEEE Trans. Image Processing **21** (2012), no. 4, 1488–1499.
8. H.-W. Chang, H. Yang, Y. Gan, and M.-H. Wang, *Sparse feature fidelity for perceptual image quality assessment*, IEEE Trans. Image Processing **22** (2013), no. 10, 4007–4018.
9. S. G. Chang, B. Yu, and M. Vetterli, *Adaptive wavelet thresholding for image denoising and compression*, IEEE Trans. Image Processing **9** (2000), no. 9, 1532–1546.
10. S. S. Channappayya, A. C. Bovik, and R. W. Heath, *A linear estimator optimized for the structural similarity index and its application to image denoising*, Proc. IEEE Int. Conf. Image Proc., IEEE, 2006, pp. 2637–2640.
11. ———, *Perceptual soft thresholding using the structural similarity index*, Proc. IEEE Int. Conf. Image Proc., IEEE, 2008, pp. 569–572.
12. S. S. Chen, D. L. Donoho, and M. A. Saunders, *Atomic decomposition by basis pursuit*, SIAM Rev. **43** (2001), 129–159.
13. G. Davis, S. Mallat, and Avellaneda. M., *Greedy adaptive approximation*, Journal of Constructive Approximation **13** (1997), 57–98.
14. D. L. Donoho, *De-noising by soft-thresholding*, IEEE Trans. Information Theory **41** (1995), no. 3, 613–627.

15. D. L. Donoho and I. M. Johnstone, *Ideal spatial adaptation by wavelet shrinkage*, Biometrika **81** (1994), no. 3, 425–455.
16. ———, *Adapting to unknown smoothness via wavelet shrinkage*, Journal of the american statistical association **90** (1995), no. 432, 1200–1224.
17. L. Jin, K. Egiazarian, and C.-C. J. Kuo, *Perceptual image quality assessment using block-based multi-metric fusion (BMMF)*, Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing, IEEE, 2012, pp. 1145–1148.
18. A. K. Katsaggelos, *Digital image restoration*, Springer Publishing Company, Incorporated, 2012.
19. A. Kolaman and O. Yadid-Pecht, *Quaternion structural similarity: a new quality index for color images*, IEEE Trans. Image Processing **21** (2012), no. 4, 1526–1536.
20. S. Mallat and Z. Zhang, *Matching pursuit with time-frequency dictionaries*, IEEE Trans. Signal Processing **41** (1993), 3397–3415.
21. J. Mannos and D. Sakrison, *The effects of a visual fidelity criterion of the encoding of images*, IEEE Trans. Information Theory **20** (1974), no. 4, 525–536.
22. K. Okarma, *Colour image quality assessment using structural similarity index and singular value decomposition*, Proc. Int. Conf. on Image Analysis and Recognition (M. Kamel and A. Campilho, eds.), LNCS, vol. 5337, Springer, Heidelberg, 2009, pp. 55–65.
23. D. Otero, D. La Torre, and E. R. Vrscay, *Structural similarity-based optimization problems with L$\hat{I}$-regularization: Smoothing using mollifiers*, Proc. Int. Conf. on Image Analysis and Recognition, Springer International Publishing, 2015, pp. 33–42.
24. D. Otero and E. R. Vrscay, *Solving optimization problems that employ structural similarity as the fidelity measure*, Proc. Int. Conf. on Image Processing, Computer Vision and Pattern Recognition, CSREA Press, 2014, pp. 474–479.
25. ———, *Unconstrained structural similarity-based optimization*, Proc. Int. Conf. on Image Analysis and Recognition, Springer International Publishing, 2014, pp. 167–176.
26. Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad, *Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition*, Proc. IEEE Asilomar Conf. on Signals, Systems, and Computers, 1993, pp. 40–44.
27. N. Ponomarenko, O. Ieremeiev, V. Lukin, K. Egiazarian, and M. Carli, *Modified image visual quality metrics for contrast change and mean shift accounting*, Proc. Int. Conf. The Experience of Designing and Application of CAD Systems in Microelectronics (2011), 305–311.
28. N. Ponomarenko, L. Jin, O. Ieremeiev, V. Lukin, K. Egiazarian, J. Astola, B. Vozel, K. Chehdi, M. Carli, F. Battisti, et al., *Image database TID2013: Peculiarities, results and perspectives*, Signal Processing: Image Communication **30** (2015), 57–77.
29. N. Ponomarenko, V. Lukin, A. Zelensky, K. Egiazarian, M. Carli, and F. Battisti, *TID2008-a database for evaluation of full-reference visual quality assessment metrics*, Advances of Modern Radioelectronics **10** (2009), no. 4, 30–45.
30. J. Portilla and E. Simoncelli, *Image restoration using gaussian scale mixtures in the wavelet domain*, Proc. IEEE Int. Conf. Image Proc., vol. 2, IEEE, 2003, pp. 965–968.
31. A. Rehman, M. Rostami, Z. Wang, D. Brunet, and E. R. Vrscay, *SSIM-inspired image restoration using sparse representation*, EURASIP Journal on Advances in Signal Processing, Special Issue on Image and Video Quality Improvement Techniques for Emerging Applications **16** (2012), no. 1, 1–12.
32. A. Rehman, Z. Wang, D. Brunet, and E. R. Vrscay, *SSIM-inspired image denoising using sparse representations*, Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (Prague, Czech Republic), 2011, pp. 1121–1124.
33. R. Rubinstein, M. Zibulevsky, and M. Elad, *Efficient implementation of the K-SVD algorithm using batch orthogonal matching pursuit*, Tech. report, Department of Computer Science, Technion, Israel Institute of Technology, Haifa, Israel, 2008.
34. M. P. Sampat, Z. Wang, S. Gupta, A. C. Bovik, and M. K. Markey, *Complex wavelet structural similarity: A new image similarity index*, IEEE Trans. Image Processing **18** (2009), no. 11, 2385–2401.
35. Z. Wang and A. C. Bovik, *A universal image quality index*, IEEE Signal Processing Letters **9** (2002), no. 3, 81–84.

36. _____, *Reduced- and no-reference image quality assessment*, IEEE Signal Processing Magazine **28** (2011), no. 6, 29 –40.
37. 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** (2004), no. 4, 600–612.
38. Z. Wang and Q. Li, *Information content weighting for perceptual image quality assessment*, IEEE Trans. Image Processing **20** (2011), no. 5, 1185–1198.
39. Z. Wang, L. Lu, and A. C. Bovik, *Video quality assessment based on structural distortion measurement*, Signal Processing: Image Communication **19** (2004), no. 2, 121–132.
40. Z. Wang and X. Shang, *Spatial pooling strategies for perceptual image quality assessment*, Proc. IEEE Int. Conf. Image Proc., IEEE, 2006, pp. 2945–2948.
41. Z. Wang, E. P. Simoncelli, and A. C. Bovik, *Multiscale structural similarity for image quality assessment*, Proc. IEEE Asilomar Conf. on Signals, Systems, and Computers, vol. 2, IEEE, 2003, pp. 1398–1402.
42. L. Zhang and H. Li, *SR-SIM: A fast and high performance IQA index based on spectral residual*, Proc. IEEE Int. Conf. Image Proc., IEEE, 2012, pp. 1473–1476.
43. L. Zhang, L. Zhang, X. Mou, and D. Zhang, *FSIM: A feature similarity index for image quality assessment*, IEEE Trans. Image Processing **20** (2011), no. 8, 2378–2386.
44. W. Zhang, A. Borji, Z. Wang, P. Le Callet, and H. Liu, *The application of visual saliency models in objective image quality assessment: A statistical evaluation*, IEEE Trans. Neural Networks and Learning Systems **27** (2016), no. 6, 1266–1278.