

Blind Deblurring Using a Simplified Sharpness Index

Arthur Leclaire, Lionel Moisan

► **To cite this version:**

Arthur Leclaire, Lionel Moisan. Blind Deblurring Using a Simplified Sharpness Index. fourth International Conference on Scale Space and Variational Methods in Computer Vision (SSVM), Jun 2013, Graz, Austria. 7893, pp.86-97, 2013, Lecture Notes in Computer Science. <10.1007/978-3-642-38267-3_8>. <hal-00795811>

HAL Id: hal-00795811

<https://hal.archives-ouvertes.fr/hal-00795811>

Submitted on 28 Feb 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Blind Deblurring Using a Simplified Sharpness Index

Arthur Leclaire and Lionel Moisan

Université Paris Descartes, France
MAP5, CNRS UMR 8145

{arthur.leclaire,lionel.moisan}@parisdescartes.fr

Abstract. It was shown recently that the phase of the Fourier Transform of an image could lead to interesting no-reference image quality measures. The Global Phase Coherence, and its recent Gaussian variant called Sharpness Index, rate the sharpness of an image in contrast not only with blur, but also noise, ringing, etc. In this work, we introduce a new variant of these indices, that can be computed with one Fourier Transform only, hence four times quicker than the Sharpness Index. We use this new index S to build an image restoration algorithm that, in a stochastic framework, selects a radial-unimodal deconvolution kernel for which the S -value of the restored image is optimal. Experiments are discussed, and comparison is made with a radial oracle deconvolution filter and the recent blind deconvolution algorithm of Levin et al.

Keywords: global phase coherence, sharpness, blind deconvolution, no-reference image quality assessment, oracle deconvolution filter

1 Introduction

No-reference image quality assessment consists in designing algorithms to evaluate the quality of an image (in particular in relation with its level of blur and noise) without requiring either an ideal version of this image (full-reference) or features extracted from this ideal image (reduced-reference). Finding good image quality (and sharpness) metrics has several applications, like, e.g., parameter selection, image restoration [16], benchmarking, or depth estimation [3].

A way to address the notion of image quality is to think in terms of precision of its geometric elements (contours, alignments, etc.). Since the pioneering work of Oppenheim and Lim [14], it is well known that the geometry of an image is mainly encoded in the phase of its Fourier Transform. And yet, the phase information itself is still very difficult to understand. A first definition of local phase coherence was given in [13], [10] and used for edge detection. Later, it was used to design a local sharpness measure in [15] and [9]. In 2008, the authors of [1] defined a notion of Global Phase Coherence (GPC), which rates the sharpness of an image depending on how the regularity of the image is destroyed as its phase information is lost. Very recently in [2], a variant of GPC called Sharpness Index (SI) was introduced. It has the advantage of being described by an explicit

closed-form formula, without needing computationally expensive Monte-Carlo Simulations like GPC.

In the present paper, we show that the SI metric can be further simplified to yield a new metric S that can be computed with only one Discrete Fourier Transform (versus four for the SI metric), while being an excellent approximation of the latter (Section 2). The behavior of this new metric is analyzed, in particular in comparison with the Q metric proposed by Zhu and Milanfar [16] (Section 3). Then, a blind deblurring algorithm is built in Section 4, that looks for the linear filter that maximizes the S -value of the restored image while imposing the Fourier Transform of the convolution kernel to be radial, unimodal and smooth. The results of this algorithm are discussed, and compared with the corresponding linear oracle and with the blind deconvolution algorithm recently proposed by Levin et al. [11].

2 Global Phase Coherence and Derived Sharpness Metrics

Let us first introduce some useful notations. In all the following, we consider gray-level images $u : \Omega \rightarrow \mathbb{R}$ defined on a discrete $M \times N$ rectangular domain

$$\Omega = \mathbb{Z}^2 \cap \left(\left[-\frac{M}{2}, \frac{M}{2} \right) \times \left[-\frac{N}{2}, \frac{N}{2} \right) \right).$$

The discrete Fourier transform (DFT) of u is the complex function \hat{u} defined by

$$\forall \boldsymbol{\xi} \in \mathbb{Z}^2, \quad \hat{u}(\boldsymbol{\xi}) = \sum_{\mathbf{x} \in \Omega} u(\mathbf{x}) e^{-i\langle \boldsymbol{\xi}, \mathbf{x} \rangle}, \quad (1)$$

where $\langle \boldsymbol{\xi}, \mathbf{x} \rangle = 2\pi \left(\frac{x_1 \xi_1}{M} + \frac{x_2 \xi_2}{N} \right)$ with $\boldsymbol{\xi} = (\xi_1, \xi_2)$ and $\mathbf{x} = (x_1, x_2)$. The function $|\hat{u}|$ will be called the modulus of u . A phase function for u is any function $\varphi : \mathbb{Z}^2 \rightarrow \mathbb{R}$ such that for all $\boldsymbol{\xi} \in \mathbb{Z}^2$, one has $\hat{u}(\boldsymbol{\xi}) = |\hat{u}(\boldsymbol{\xi})| e^{i\varphi(\boldsymbol{\xi})}$.

The Ω -periodization of u is the image $\dot{u} : \mathbb{Z}^2 \rightarrow \mathbb{R}$ that extends u to \mathbb{Z}^2 by $\dot{u}(\mathbf{x}) = u(\mathbf{x}')$, where \mathbf{x}' is the unique element of Ω such that $\mathbf{x}' - \mathbf{x} \in M\mathbb{Z} \times N\mathbb{Z}$. The gradient of \dot{u} is defined by

$$\forall (x, y) \in \mathbb{Z}^2, \quad \nabla \dot{u}(x, y) = \begin{pmatrix} \partial_x \dot{u}(x, y) \\ \partial_y \dot{u}(x, y) \end{pmatrix} = \begin{pmatrix} \dot{u}(x+1, y) - \dot{u}(x, y) \\ \dot{u}(x, y+1) - \dot{u}(x, y) \end{pmatrix}, \quad (2)$$

and the (periodic and anisotropic) Total Variation of u is

$$\text{TV}(u) = \|\partial_x \dot{u}\|_1 + \|\partial_y \dot{u}\|_1 = \sum_{\mathbf{x} \in \Omega} |\partial_x \dot{u}(\mathbf{x})| + |\partial_y \dot{u}(\mathbf{x})|. \quad (3)$$

The autocorrelation of $\nabla \dot{u}$ is the function $\Gamma : \Omega \rightarrow \mathbb{R}^{2 \times 2}$ defined by

$$\Gamma(\mathbf{z}) = \begin{pmatrix} \Gamma_{xx}(\mathbf{z}) & \Gamma_{xy}(\mathbf{z}) \\ \Gamma_{xy}(\mathbf{z}) & \Gamma_{yy}(\mathbf{z}) \end{pmatrix} = \sum_{\mathbf{y} \in \Omega} (\nabla \dot{u}(\mathbf{y})) (\nabla \dot{u}(\mathbf{y} + \mathbf{z}))^T. \quad (4)$$

2.1 Global Phase Coherence

As mentioned earlier, the phase of an image u encodes a great part of the geometry of u : if one reproduces the famous experiment of [14] consisting in imposing the phase of an image u to another image v , one can see on the result that several edges from u have appeared and all the geometric content of v has disappeared. Indeed, phase coefficients need strong alignment constraints in order to produce sharp edges and clean flat regions in an image.

The Global Phase Coherence metric introduced in [1] quantifies how the loss of this phase coherence affects the image regularity, measured by its Total Variation (3). More precisely, the phase of an image u is randomized to produce the Random Phase Noise image U_ψ defined in Fourier Domain by

$$\forall \boldsymbol{\xi} \in \Omega, \quad \widehat{U}_\psi(\boldsymbol{\xi}) = |\widehat{u}(\boldsymbol{\xi})| e^{i\psi(\boldsymbol{\xi})}, \quad (5)$$

where $\psi : \Omega \rightarrow \mathbb{R}$ is a uniform random phase (the coefficients of ψ are independent and uniformly distributed in $(0, 2\pi)$, modulo the relation $\psi(-\boldsymbol{\xi}) = -\psi(\boldsymbol{\xi})$ ensuring that U_ψ is real-valued, see [8]), which leads to

Definition 1 (Blanchet, Moisan, Rougé, 2008 [1]). *The Global Phase Coherence (GPC) of u is the number*

$$\text{GPC}(u) = -\log_{10} \mathbb{P}(\text{TV}(U_\psi) \leq \text{TV}(u)). \quad (6)$$

For an image u with sharp edges and clean uniform zones, the Total Variation is expected to be low amongst the ones of its phase randomizations. Therefore, for such an image, the probability of the event $\{\text{TV}(U_\psi) \leq \text{TV}(u)\}$ will be very small, and the value of $\text{GPC}(u)$ will be large. That is why this phase coherence index (and the variants that follow) is expected to behave like an image quality measure. Note that without the logarithm in (6), the values of $\text{GPC}(u)$ would often cause a numerical underflow (a value like, e.g., 10^{-1000} cannot be represented in most computer environments).

The main issue with (6) is that no closed-form formula has been found so far to compute $\text{GPC}(u)$, so that a computationally expensive Monte-Carlo simulation (coupled with a Gaussian approximation of the random variable $\text{TV}(U_\psi)$) is proposed in [1], which limits the potential application of the GPC metric.

2.2 Sharpness Index

Hopefully, a closed-form variant of GPC was recently found. In [2], the periodic convolution of u with a conveniently normalized Gaussian white noise W is considered instead of U_ψ , and the first two moments of $\text{TV}(u * W)$ are explicitly computed in function of u . Note that $u * W$ is nothing but the natural Gaussian approximation of U_ψ , and these two random images only differ by a convolution with the texture of a white noise, which is close to a Dirac distribution [6] (in Fourier Domain, $\widehat{u * W}$ and \widehat{U}_ψ differ by a multiplicative Rayleigh Noise). Even if the exact law of $\text{TV}(u * W)$ seems difficult to compute, it is expected to be approximately Gaussian, which leads to

Definition 2 (Blanchet, Moisan, 2012 [2]). *The Sharpness Index of u is*

$$\text{SI}(u) = -\log_{10} \Phi \left(\frac{\mu - \text{TV}(u)}{\sigma} \right) \quad (7)$$

$$\text{where } \Phi(t) = \frac{1}{\sqrt{2\pi}} \int_t^{+\infty} \exp\left(-\frac{x^2}{2}\right) dx, \quad \mu = (\alpha_x + \alpha_y) \sqrt{\frac{2}{\pi}} \sqrt{MN}, \quad (8)$$

$$\sigma^2 = \frac{2}{\pi} \sum_{\mathbf{z} \in \Omega} \alpha_x^2 \cdot \omega \left(\frac{\Gamma_{xx}(\mathbf{z})}{\alpha_x^2} \right) + 2\alpha_x \alpha_y \cdot \omega \left(\frac{\Gamma_{xy}(\mathbf{z})}{\alpha_x \alpha_y} \right) + \alpha_y^2 \cdot \omega \left(\frac{\Gamma_{yy}(\mathbf{z})}{\alpha_y^2} \right), \quad (9)$$

$\alpha_x = \|\partial_x \dot{u}\|_2 = (\sum_{\mathbf{x} \in \Omega} |\partial_x \dot{u}(\mathbf{x})|^2)^{\frac{1}{2}}$, $\alpha_y = \|\partial_y \dot{u}\|_2$, Γ is the autocorrelation of $\nabla \dot{u}$ given in (4), and ω is the function defined by

$$\forall t \in [-1, 1], \quad \omega(t) = t \cdot \text{Arcsin}(t) + \sqrt{1 - t^2} - 1. \quad (10)$$

In practice, the numerical computation of $\text{SI}(u)$ requires the computation of $\text{TV}(u)$, α_x and α_y (linear time), plus the three different components of Γ that can be computed with four Fast Fourier Transforms (a direct FFT of u , and 3 inverse FFTs for the cross correlations of $\partial_x \dot{u}$ and $\partial_y \dot{u}$). The overall cost is hence dominated by these four $M \times N$ FFT computations, which represents a complexity of $\mathcal{O}(MN \log(MN))$ for well-suited image dimensions. Note that the function Φ is available in most mathematical libraries through the complementary error function (often written `erfc`), but when t is greater than say, 20, the following (almost exact) approximation is systematically used to avoid numerical underflow:

$$-\log_{10} \Phi(t) \simeq \frac{t^2 + \log(2\pi t^2)}{2 \log(10)}. \quad (11)$$

2.3 A Simplified Version of SI

We now introduce a new index S which is analytically close to SI and faster to compute. For that, let us observe that $\omega(t)$ is equivalent to $t^2/2$ when $t \rightarrow 0$. Therefore, this approximation can be used to replace $\sigma^2 = \text{Var}(\text{TV}(u * W))$ by

$$\sigma_a^2 = \frac{1}{\pi} \sum_{\mathbf{z} \in \Omega} \alpha_x^2 \cdot \left(\frac{\Gamma_{xx}(\mathbf{z})}{\alpha_x^2} \right)^2 + 2\alpha_x \alpha_y \cdot \left(\frac{\Gamma_{xy}(\mathbf{z})}{\alpha_x \alpha_y} \right)^2 + \alpha_y^2 \cdot \left(\frac{\Gamma_{yy}(\mathbf{z})}{\alpha_y^2} \right)^2,$$

which after simplification yields the following

Definition 3. *The S -metric of an image u is*

$$S(u) = -\log_{10} \Phi \left(\frac{\mu - \text{TV}(u)}{\sigma_a} \right), \quad (12)$$

$$\text{where } \sigma_a^2 = \frac{1}{\pi} \left(\frac{\|\Gamma_{xx}\|_2^2}{\alpha_x^2} + 2 \cdot \frac{\|\Gamma_{xy}\|_2^2}{\alpha_x \alpha_y} + \frac{\|\Gamma_{yy}\|_2^2}{\alpha_y^2} \right), \quad (13)$$

and α_x , α_y , μ , and Γ are as in Definition 2.

Whereas SI needed all the coefficients of the gradient auto-correlation matrix, the S metric only depends on the overall energy of the three components Γ_{xx} , Γ_{xy} and Γ_{yy} . Thanks to Parseval's formula, they can be computed in Fourier domain, recalling that

$$\widehat{\Gamma_{xx}}(\boldsymbol{\xi}) = |\widehat{\partial_x \dot{u}}(\boldsymbol{\xi})|^2 = 2 \sin^2 \left(\frac{\pi \xi_1}{M} \right) |\hat{u}(\boldsymbol{\xi})|^2, \quad \widehat{\Gamma_{yy}}(\boldsymbol{\xi}) = 2 \sin^2 \left(\frac{\pi \xi_2}{N} \right) |\hat{u}(\boldsymbol{\xi})|^2,$$

$$\text{and } |\widehat{\Gamma_{xy}}(\boldsymbol{\xi})| = |\widehat{\partial_x \dot{u}}(\boldsymbol{\xi})| |\widehat{\partial_y \dot{u}}(\boldsymbol{\xi})| = 2 \left| \sin \left(\frac{\pi \xi_1}{M} \right) \sin \left(\frac{\pi \xi_2}{N} \right) \right| |\hat{u}(\boldsymbol{\xi})|^2.$$

The computation of S only involves the l^1 and l^2 norms of the gradient, and the FFT of u . Thus, the overall dominant cost is only one FFT (compared to 4 for SI), while the approximation of SI by S is very good as stated by

Proposition 1. *We have $0 \leq \frac{\sigma^2 - \sigma_a^2}{\sigma_a^2} \leq \pi - 3 \approx 0.142$.*

Proof. With the expressions of σ^2 and σ_a^2 , one can write

$$\sigma^2 - \sigma_a^2 = \frac{2}{\pi} \sum_{\mathbf{x} \in \Omega} \alpha_x^2 \left[\omega \left(\frac{\Gamma_{xx}(\mathbf{x})}{\alpha_x^2} \right) - \frac{1}{2} \left(\frac{\Gamma_{xx}(\mathbf{x})}{\alpha_x^2} \right)^2 \right]$$

$$+ 2\alpha_x \alpha_y \left[\omega \left(\frac{\Gamma_{xy}(\mathbf{x})}{\alpha_x \alpha_y} \right) - \frac{1}{2} \left(\frac{\Gamma_{xy}(\mathbf{x})}{\alpha_x \alpha_y} \right)^2 \right] + \alpha_y^2 \left[\omega \left(\frac{\Gamma_{yy}(\mathbf{x})}{\alpha_y^2} \right) - \frac{1}{2} \left(\frac{\Gamma_{yy}(\mathbf{x})}{\alpha_y^2} \right)^2 \right].$$

Besides, Taylor Formula applied to ω yields

$$\forall t \in [-1, 1], \quad 0 \leq \omega(t) - \frac{1}{2}t^2 \leq ct^4 \leq ct^2, \quad (14)$$

with $c = \omega(1) - 1 = \frac{\pi-3}{2}$, and thus

$$0 \leq \sigma^2 - \sigma_a^2 \leq \frac{2c}{\pi} \sum_{\mathbf{x} \in \Omega} \alpha_x^2 \left(\frac{\Gamma_{xx}(\mathbf{x})}{\alpha_x^2} \right)^2 + 2\alpha_x \alpha_y \left(\frac{\Gamma_{xy}(\mathbf{x})}{\alpha_x \alpha_y} \right)^2 + \alpha_y^2 \left(\frac{\Gamma_{yy}(\mathbf{x})}{\alpha_y^2} \right)^2 = 2c\sigma_a^2.$$

3 Validation of S as a Quality Measure

As in [1], we shall systematically apply two simple image transforms to an image u before computing $S(u)$ with (12), in order to avoid periodization and quantization biases. First, as the S metric is defined (like GPC and SI) through a periodic setting, the periodic component of u (see [12]) is first extracted to avoid discontinuities across the image frame border. Then, a simple dequantization procedure (a $(1/2, 1/2)$ sub-pixel translation with Fourier interpolation, see [5]) is applied to ensure that the quantization of the original image (generally with 256 gray levels) does not artificially decrease its Total Variation.

On Fig. 1, we give a first empirical evidence that S behaves as an image quality measure. Indeed, this blur-noise diagram shows that the value of S decreases




Fig. 1. Some level lines of the function $(r, \beta) \mapsto S(g_r * u + \beta W)$ where g_r is the 2-D Gaussian convolution kernel with standard deviation r , and W is a white noise image with unit variance in each pixel. The absolute values of S and the exact shape of the level lines depend on the image considered (here, Barbara), but the overall shape remains similar.

as the the level of blur or noise increases, with a correspondence between noise and blur which is similar to GPC and SI (see [1, 2]).

The S metric not only decreases with respect to blur and noise, but unlike the Q metric of Zhu and Milanfar [16], it also decreases when ringing artifacts (that may result from excessive deblurring) appear, as shown on Fig. 2. This suggests that S could be used in a parametric or even non-parametric blind deblurring algorithm, as will be done in Section 4. Note however, that the Q metric performs slightly better than S for parameter selection in the SKR denoising method presented in [16]. In fact, the reason is logical considering the origin of S : the noise left by SKR in uniform zones is structured and the coherence of its phase makes S prefer less denoised images than Q .

4 Application to Blind Deblurring

Removing blur from a single image is a difficult task. If the blur is linear and spatially uniform, it can be modeled as a convolution. Several algorithms (see, e.g., the recent efficient scheme for $TV - L^2$ deblurring proposed in [4]), have been proposed to invert the effect of this convolution when the blur kernel is known. Addressing the problem of blind deconvolution, i.e. when the kernel is not known, is even more difficult, and several solutions have been proposed in the last decades. In the present paper, we shall consider in particular (for comparison purposes) the very recent work of Levin et al. [11].

Here, rather than trying to reverse the effect of a convolution, we shall try to improve the image directly by convolving the blurry image u with a unit-mass (that is, average-preserving) filter k that maximizes $F_u(k) = S(k * u)$. Indeed, Fig. 2 shows that, in a parametric Wiener deconvolution, the S metric is able to select the blur parameter. In [3], Calderero and Moreno made a similar observation for the SI metric in a context of reverse diffusion. In this section, we will show that S can be used for non-parametric blind deblurring.




Fig. 2. These diagrams plot the proposed metric S (left) and the metric Q (right) of Zhu and Milanfar [16] obtained when applying (in Fourier domain) a H^1 -regularized deconvolution filter to a natural image (Yale). The abscissa (ρ) corresponds to the standard deviation of the supposedly Gaussian convolution kernel. Over the results obtained for the whole range of ρ , one can see that contrary to Q , S attains a maximum for $\rho = 0.55$, which roughly corresponds to the value beyond which ringing artifacts begin to appear. This suggests that S is able to discriminate ringing, whereas Q does not.

4.1 Kernels with Compact Support

Algorithm 1 below can be used to optimize F_u on particular sets of kernels k , as soon as the number of coefficients that define k remains small enough. For example, one can optimize F_u on the set of symmetric 5×5 kernels (12 coefficients), or on the set of separable symmetric 21×21 kernels (20 coefficients, or 10 if the same one-dimensional kernel is used for each coordinate). In general, the results obtained with Algorithm 1 are good, but some images lead to interesting failure cases, in particular when regions with highly structured textures or dominant orientations are present (see Fig. 3). Since the functional F_u is not concave, it does not necessarily have a unique local maximum, and a reason could be that Algorithm 1 does not manage to converge to the actual global maximum of the objective function F_u . However, experiments suggest that the failure is more likely to be due to an inadequate set of kernels (in particular the set of separable kernels). To avoid such degenerated cases, and get rid of the small-kernel-support constraint, we consider in Section 4.2 other sets of kernels for which constraints are considered in Fourier domain.

Algorithm 1

- Begin with $k = \delta_0$ (discrete Dirac kernel)
- Repeat n times
 - ▷ Define k' from a random perturbation of k
 - ▷ If $S(k' * u) > S(k * u)$ then $k \leftarrow k'$
- Renormalize k to a unit-mass kernel
- Return k and $k * u$




Fig. 3. Blind deblurring of the original *Room* image (left). Algorithm 1 is applied for two different sets of kernels: 5×5 symmetric kernels (second column) and 21×21 symmetric separable kernels (third column). The right image is the result obtained with the method proposed in Section 4.2 (Algorithm 2, $\mu = 0$). We observe that Algorithm 1 fails in both cases (probably in reason of the large striped texture), while the radial-unimodal constraint imposed in Algorithm 2 yields a nice-looking result.

4.2 Optimization of a Radial-Unimodal Kernel

Instead of imposing that the convolution kernel has a fixed compact support, we here consider the set of kernels that are radial in Fourier domain, with a unimodal profile, which is a plausible assumption for a deconvolution kernel. More precisely, we assume that the DFT of the restoration kernel k_r is given by

$$\forall \xi \in \Omega, \quad \widehat{k}_r(\xi_1, \xi_2) = L_r \left(\sqrt{2(d-1) \left(\left(\frac{\xi_1}{M} \right)^2 + \left(\frac{\xi_2}{N} \right)^2 \right)} \right), \quad (15)$$

where $L_r : [0, d-1] \rightarrow \mathbb{R}$ is the piecewise affine interpolate on $[0, d-1]$ of the finite sequence $r(0) = 1, r(1), r(2), \dots, r(d-2), r(d-1) = 0$. This sequence is supposed to be unimodal, which means that there exists a value i_m (mode index) that satisfies

$$\forall i < i_m, \quad r(i+1) \geq r(i), \quad \text{and} \quad \forall i \geq i_m, \quad r(i+1) \leq r(i).$$

One possible perturbation strategy for this set of kernels consists in the addition of a uniform random value to a randomly chosen coefficient of r , followed by a projection on the set U of unimodal sequences (this projection can be computed in $O(n^2)$ operations with the Pool Adjacent Violators algorithm [7]). We observed that with this strategy, moving the mode position was difficult, so we relaxed the unimodality constraint and incorporated in the objective function the l^2 -distance $d(r, U)$ between r and the set U of unimodal sequences. We also found useful to add the possibility to increase the regularity of the radial profile r by incorporating a term depending on

$$\|r\|_{H^1}^2 = \sum_{i=0}^{d-2} (r(i+1) - r(i))^2. \quad (16)$$

Finally, the objective function (to be maximized) is

$$\mathcal{F}_u(r) = S(k_r * u) - \lambda d(r, U) - \mu \|r\|_{H^1}, \quad (17)$$

where λ and μ are weighting parameters.

In order to maximize \mathcal{F}_u , we used Algorithm 2 below. We observed that $n = 10000$ was sufficient to ensure convergence on r (the relative changes after 10000 iterations were less than 10^{-3}); moreover, we checked that several realizations of this stochastic technique led to the same local maximum. The other parameters were set to $d = 20$, initial $i_m = 5$, $a = 0.1$, $\lambda = 10000$, and μ varying from 0 to 100. Let us comment the choice of i_m . As soon as several local maxima are present, the result of an optimization technique may depend on the initialization, and the natural solution would be to apply the algorithm for all possible values of the initial mode index. But in a wide majority of cases, we observed that the result of this algorithm was not depending on this initial value. In a few cases, two different local maxima could be found, but the higher value of the objective function \mathcal{F}_u was always obtained for $i_m \in [d/4, 3d/4]$. This is why it seems empirically sufficient to run the algorithm only once with the initial value of i_m in that range.

Algorithm 2

- Initialize r with the piecewise-linear profile such that $r(0) = 1$, $r(i_m) = 2$, and $r(d-1) = 0$.
- Repeat n times
 - ▷ Pick a random index $i \in [1, d-2]$
 - ▷ Draw a random value $\varepsilon \in [-a/2, a/2]$ with uniform distribution
 - ▷ Set $r' \leftarrow r$, then $r'(i) \leftarrow r(i) + \varepsilon$
 - ▷ If $\mathcal{F}_u(r') > \mathcal{F}_u(r)$ then $r \leftarrow r'$
- Return r , k_r and $k_r * u$

4.3 Results

We first used Algorithm 2 with $\mu = 0$ on *Room* image, and checked that the failure of Algorithm 1 was avoided (Fig. 3, right).

Then, to produce the results shown in Fig. 4, we took two classical images (*Capitol* and *Parrots*) and corrupted them with a Gaussian blur kernel (width 1 pixel) and an additive white Gaussian noise (variance 1). We then applied several deblurring algorithms (detailed below) and evaluated their performances by computing their respective PSNR values with respect to the original clean image. Notice, however, that for blind deblurring tasks the PSNR value is not very reliable (in particular because even the original clean image is necessarily, in some sense, blurry and noisy), and visual inspection is often preferable to compare the algorithms.

First, we used Algorithm 2 with $\mu = 0$ and $\mu = 10$. We observed that the results were stable, and that the restoration resulted in a significant sharpness increase. However, for $\mu = 0$ some low-frequency noise is still visible on uniform zones. Increasing the value of μ to $\mu = 10$ reduces the residual noise (because it reduces the amplitude of L_r , that is, the amplification of noisy Fourier coefficients), but also attenuate some details in textured zones.

We then compared these results with the state-of-the-art blind deconvolution algorithm of Levin et al. [11]. One can see in Fig. 4 that Algorithm 2 is more precise on the fine details of the image, but it also keeps much more noise than this method. In fact, the result of [11] is really “clean” and has a small Total Variation, which explains incidentally why its S value is significantly larger compared to Algorithm 2. Notice also that the method [11] is more general, and has been shown to perform particularly well in the case of a motion blur, while the radial constraint of Algorithm 2 cannot handle such motion blurs.

Another interesting experiment consists in computing the optimal kernel k_o that maximizes the expected distance between the reconstructed image and the clean image u_0 , knowing the parameters (kernel κ and noise level β) of the degradation process (this is an oracle since neither these parameters nor the clean image are supposed to be known). One has

$$k_o = \underset{k}{\text{Arg min}} \mathbb{E} \|u_0 - k * (\kappa * u_0 + \beta W)\|^2 \quad (18)$$

and if all kernels k were considered, the solution would be given by

$$\forall \xi \in \Omega, \quad \widehat{k}_o(\xi) = \frac{\overline{\widehat{\kappa}(\xi)} |\widehat{u}_0(\xi)|^2}{|\widehat{\kappa}(\xi)|^2 |\widehat{u}_0(\xi)|^2 + \sigma^2 MN}. \quad (19)$$

Now, since we only consider kernels whose DFT is built from a radial profile (linearly interpolated on d points), one can show that the optimal radial profile is the minimum of a quadratic function and thus can be obtained by solving a small linear system. In Fig. 5 we can see that the profile of the oracle radial kernel is unimodal, with a mode at a position which is close to the one estimated by Algorithm 2. The restored images obtained with this oracle filter are also displayed in Fig. 4: they are a little more precise on the details, but present a significant amount of structured noise; indeed, such a noise is not very costly for a l^2 risk function. This is somehow reassuring: this shows that the structured noise also appearing with Algorithm 2 (in particular with $\mu = 0$) is truly a limit of techniques based on linear filtering.

5 Conclusion

We introduced a new variant of the GPC and SI image quality metrics, that can be computed four times faster than SI. This new index S provides a sharpness measure that can be used in a stochastic optimization framework to achieve blind deblurring through linear convolution with a radial unimodal kernel. Though suffering from the limits of linear filtering, the obtained results are convincing, and visually similar to the best possible ones (oracle) obtained by such an approach. The extension to motion blur kernels could be an interesting generalization, as well as the use of more sophisticated (non-linear) restoration techniques to make the best usage of the selection performances of the S metric.

Acknowledgments. This work has been supported by the French National Research Agency under grant ANR-09-BLAN-0029-01.




Fig. 4. Blind deblurring of a degraded version of *Capitol* and *Parrots* images (Gaussian blur of width 1 pixel plus Gaussian noise with variance 1). We present in each case the original image, the blurred and noisy input, the result of Levin et al. algorithm [11], the oracle output (best possible result obtained by a radial unimodal convolution filter) and the results of Algorithm 2 with $\mu = 0$ and $\mu = 10$. The PSNR values are computed in each case with respect to the original image.




Fig. 5. Different radial profiles (oracle, and Algorithm 2 with $\mu = 0$ and $\mu = 10$) obtained on images *Capitol* (left) and *Parrots* (right).

References

1. Blanchet, G., Moisan, L., Rougé, B.: Measuring the Global Phase Coherence of an Image. In Proceedings of ICIP'08, pp. 1176–1179 (2008).
2. Blanchet, G., Moisan, L.: An Explicit Sharpness Index Related to Global Phase Coherence. In Proceedings of ICASSP'12, pp. 1065–1068 (2012).
3. Calderero, F., Moreno, P.: Evaluation of Sharpness Measures and Proposal of a Stop Criterion for Reverse Diffusion in the Context of Image Deblurring. In Proceedings of VISAPP'13 (2013).
4. Chambolle, A., Pock, T.: A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging. *Journal of Mathematical Imaging and Vision*, vol. 40:1, pp. 161–179 (2011).
5. Desolneux, A., Moisan, L., Morel, J.-M.: Dequantizing image orientation. *IEEE Transactions on Image Processing*, vol. 11:10, pp. 1129–1140 (2002).
6. Desolneux, A., Moisan, L., Ronsin, S.: A Compact Representation of Random Phase and Gaussian Textures. In Proceedings of ICASSP'12, pp. 1381–1384 (2012).
7. Frisen, M.: Unimodal Regression. *Journal of the Royal Statistical Society. Series D The Statistician*, vol. 35:4, pp. 479–385 (1986).
8. Galerne, B., Gousseau, Y., Morel, J.-M.: Random Phase Textures : Theory and Synthesis. *IEEE Transactions on image processing*, vol. 20:1, pp. 257–267 (2011).
9. Hassen, R., Wang, Z., Salama, M.: No-Reference Image Sharpness Assessment Based on Local Phase Coherence Measurement. *Proc. of ICASSP'10*, pp. 2434–2437 (2010).
10. Kovési, P.: Phase Congruency: a Low-level Image Invariant. *Psychological Research*, vol. 64, pp. 136–148 (2000).
11. Levin, A., Weiss, Y., Durand, F., Freeman, W.T.: Efficient Marginal Likelihood Optimization in Blind Deconvolution. In Proceedings of CVPR'11 (2011).
12. Moisan, L.: Periodic Plus Smooth Image Decomposition. *Journal of Mathematical Imaging and Vision*, vol. 39:2, pp. 120–145 (2011).
13. Morrone, M.C., Burr, D.C.: Feature Detection in Human Vision: a Phase-Dependent Energy Model. *Proc. R. Soc. Lond. B235*, pp. 221–245 (1988).
14. Oppenheim, A.V., Lim, J.S.: The Importance of Phase in Signals. *Proceedings of the IEEE*, vol. 69, pp. 529–541 (1981).
15. Wang, Z., Simoncelli, E.P.: Local Phase Coherence and the Perception of Blur. In *Adv. Neural Information Processing Systems (NIPS03)*, vol. 16, pp. 786–792 (2004).
16. Zhu, X., Milanfar, P.: Automatic Parameter Selection for Denoising Algorithms Using a No-Reference Measure of Image Content. *IEEE Transactions on Image Processing*, vol. 19:12, pp. 3116–3132 (2010).