

A NO-REFERENCE SHARPNESS METRIC SENSITIVE TO BLUR AND NOISE

Xiang Zhu and Peyman Milanfar

Electrical Engineering Department
University of California at Santa Cruz, CA, 95064
xzhu@soe.ucsc.edu

ABSTRACT

A no-reference objective sharpness metric detecting both blur and noise is proposed in this paper. This metric is based on the local gradients of the image and does not require any edge detection. Its value drops either when the test image becomes blurred or corrupted by random noise. It can be thought of as an indicator of the signal to noise ratio of the image. Experiments using synthetic, natural, and compressed images are presented to demonstrate the effectiveness and robustness of this metric. Its statistical properties are also provided.

Index Terms— Sharpness metric, blur, noise, singular value, gradient, covariance.

1. INTRODUCTION

A no-reference image quality metric can play a significant role in image processing, as it can be used in many applications such as algorithm performance analysis, algorithm parameter optimization, video compression and communication. Various methods focusing on sharpness or blur measurement have been proposed. Some of them are based on pixel derivatives [1], kurtosis [2], and DCT coefficients [3], and a large group of metrics analyze the spread of the edges in an image [4], [5]. Although these metrics are often applied to assess the image quality, most neglect the effects from other possible degradation sources, such as random noise.

Some sharpness metrics perform well in detecting blur in the presence of white Gaussian noise (WGN) [6], [7]; an example is the Riemannian tensor based metric [6], whose value drops when the image is increasingly more blurred. However, the value of this measure rises if the variance of noise is increased, which means that it can not be used to distinguish image quality decay against high frequency behavior due to noise.

In this paper we propose a new sharpness metric that can quantify the amount of both blur and random noise. The value of this metric drops either when image becomes blurred or noisy, so that it can be used to capture the change of visual quality in many image processing applications. The metric is

based on the singular value decomposition of the local image gradient matrix. We first introduce the gradient matrix and its largest singular value s_1 for various types of local image regions, and characterize the behavior of s_1 in the presence of blur and noise. Then we give the definition of our sharpness metric H . Several experiments are provided to highlight the performance of the metric. Conclusions and further research directions based on metric H are provided in the final section.

2. SINGULAR VALUES OF THE IMAGE GRADIENT MATRIX

Image structure can be measured effectively by using the differences in pixels (or image gradients). Consider an image of interest $g(x, y)$. The gradient matrix of a region within an $N \times N$ local analysis window (w_i) is defined as:

$$\mathbf{G} = \begin{bmatrix} \vdots & \vdots \\ g_x(k) & g_y(k) \\ \vdots & \vdots \end{bmatrix}, \quad k \in w_i \quad (1)$$

where $[g_x(k), g_y(k)]$ denote the gradient of the image at point (x_k, y_k) . The corresponding gradient covariance matrix is

$$\mathbf{C} = \mathbf{G}^T \mathbf{G} = \begin{bmatrix} \sum_{k \in w_i} g_x^2(k) & \sum_{k \in w_i} g_x(k)g_y(k) \\ \sum_{k \in w_i} g_x(k)g_y(k) & \sum_{k \in w_i} g_y^2(k) \end{bmatrix} \quad (2)$$

Not surprisingly, some interesting information about the content of the image patch w_i can be derived from the gradient matrix \mathbf{G} or the gradient covariance matrix \mathbf{C} . One example is to calculate the local dominant orientation by computing the (compact) Singular Value Decomposition (SVD) of \mathbf{G} [8]

$$\mathbf{G} = \mathbf{U} \mathbf{S} \mathbf{V}^T = \mathbf{U} \begin{bmatrix} s_1 & 0 \\ 0 & s_2 \end{bmatrix} [\mathbf{v}_1 \quad \mathbf{v}_2]^T \quad (3)$$

where both \mathbf{U} and \mathbf{V} are orthonormal matrices. The column vector \mathbf{v}_1 represents the dominant orientation of the local gradient field. Correspondingly, the second singular vector \mathbf{v}_2 (which is orthogonal to \mathbf{v}_1) will describe the dominant orientation of this patch. The singular values $s_1 \geq s_2 \geq 0$ represent the energy in the directions \mathbf{v}_1 and \mathbf{v}_2 , respectively.

This work was supported in part by AFOSR grant FA9550-07-1-0365.




Fig. 1. Types of patches that are used in the experiments through this paper.

The above quantities can equivalently be measured using the eigenvectors of \mathbf{C} , because

$$\mathbf{C} = \mathbf{V}\mathbf{S}^T\mathbf{S}\mathbf{V}^T = \mathbf{V} \begin{bmatrix} s_1^2 & 0 \\ 0 & s_2^2 \end{bmatrix} \mathbf{V}^T. \quad (4)$$

As we will describe below, since the singular values reflect the strength of the gradients along dominant direction and its perpendicular direction, they are sensitive to blurring, and therefore, may be used to define a sharpness metric. Before we proceed with the definition of our sharpness metric, we analyze the behavior of s_1 and s_2 on several types of idealized patches. They include flat, linear, quadratic, and edged patches (shown in Fig. 1).

In the flat case, all points within the $N \times N$ patch share a common intensity value:

$$g(x_k, y_k) = c \quad (5)$$

Both $g_x(k)$ and $g_y(k)$ are equal to 0 for $k = 1, 2, \dots, N^2$, and $s_1 = s_2 = 0$. Naturally, ignoring boundary effect arising from the finite nature of the window, a flat patch remains unchanged after being blurred. In what follows, we will apply a space invariant Gaussian blur function iteratively to the canonical regions shown in Fig. 1, and observe how the singular values behave. In this sense, the flat region can be thought of as the "final" result as the number of iterations (or equivalently the strength of the blur) is made arbitrarily large.

In the linear patch, the gray value of each point can be modeled as:

$$g(x_k, y_k) = a(x_k \cos \theta + y_k \sin \theta) + b \quad (6)$$

where a decides the slope, θ decides the orientation, and b is the bias. It can be deduced that s_1 and s_2 have the following values:

$$\begin{aligned} s_1 &= aN \\ s_2 &= 0 \end{aligned} \quad (7)$$

Both s_1 and s_2 are independent from the orientation, and s_1 is proportional to the slope given a fixed patch size, while s_2 remains at zero.

For the sake of simplicity, the quadratic patch is modeled as:

$$g(x_k, y_k) = a_1(x_k - x_c)^2 + a_2(y_k - y_c)^2 \quad (8)$$

where (x_c, y_c) is the center point. The singular values of its gradient matrix are:

$$\begin{aligned} s_1 &= a_{max}N\sqrt{\frac{(N-1)(N+1)}{3}} \\ s_2 &= a_{min}N\sqrt{\frac{(N-1)(N+1)}{3}} \end{aligned} \quad (9)$$

where

$$\begin{aligned} a_{max} &= \max(a_1, a_2) \\ a_{min} &= \min(a_1, a_2) \end{aligned} \quad (10)$$

Here s_1 and s_2 reflect the value of a_1 and a_2 , which determine the slope at each point, and thus determine the sharpness of the region.

Another type of image region that is very sensitive to blurring is the ideal edged patch. Again, in the interest of convenience we just look at an ideal vertical edge:

$$g(x_k, y_k) = \begin{cases} b + c & x_n > x_c \\ b & otherwise \end{cases} \quad (11)$$

where, without loss of generality, c is a positive value. The corresponding singular values are:

$$\begin{aligned} s_1 &= \frac{c}{2}\sqrt{2N} \\ s_2 &= 0 \end{aligned} \quad (12)$$

Only s_1 here reflects the value of parameter c , which gives the intensity difference between two sides of the edge.

In general, regardless of the patch type, rotating an image patch by an arbitrary angle θ will not change the singular values of the gradient matrix. The relationship between the rotated gradient matrix \mathbf{G}_θ and the unrotated \mathbf{G} is

$$\mathbf{G}_\theta = \mathbf{G}\mathbf{R}_\theta^T \quad (13)$$

where \mathbf{R}_θ is the (orthonormal) rotation matrix:

$$\mathbf{R}_\theta = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \quad (14)$$

Therefore, the SVD of \mathbf{G}_θ becomes:

$$\mathbf{G}_\theta = \mathbf{U}\mathbf{S}(\mathbf{R}_\theta\mathbf{V})^T \quad (15)$$

This equation says that the directions \mathbf{v}_1 and \mathbf{v}_2 are correspondingly rotated, but the singular values s_1 and s_2 remain unchanged.

It is observed through the above analysis that the singular value s_1 is quite intimately related with the sharpness of the local region. This is valid not only in regions with strong direction and contrast (edged patch), but also in regions which may be isotropic (quadratic patch, where $a_1 = a_2$), or very smooth (linear patch).

To verify the usefulness of s_1 in the presence of blur, we applied a Gaussian blur (of size 5×5 and variance 0.25) iteratively to the above patch types and recorded the resulting s_1 values, which are shown in Fig. 2. The size of the patches is




Fig. 2. Plots of s_1 in blurring process for different patches. Patch size $N = 7$. The distribution of the smoothing kernel is Gaussian, and its standard deviation $\sigma_G = 0.5$. We apply this kernel iteratively to make each patch more and more blurred.

11×11 , and we only analyze the 7×7 window in the center to avoid border effects. It is observed that as the blurring iteration goes on, s_1 for all the non-flat patches drops steadily as expected.

Next, we take noise into account. A good sharpness metric should react reasonably to both blur and random noise. So the next question we address is what happens to s_1 if the image (or patch) is corrupted by white Gaussian noise (WGN).

Assume that we have an $N \times N$ image patch denoted in column-stacked vector format as an $N^2 \times 1$ vector \mathbf{n} , which contains i.i.d. samples of zero-mean WGN with variance σ^2 . In practice the statistics of \mathbf{G}_n depend upon the way we calculate the discrete derivatives. For example, the gradient of \mathbf{n} in x and y directions can be produced by applying the filters:

$$\frac{1}{2} \times \begin{bmatrix} 0 & 0 & 0 \\ -1 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}, \quad \frac{1}{2} \times \begin{bmatrix} 0 & -1 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \quad (16)$$

or the filter masks:

$$\frac{1}{8} \times \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}, \quad \frac{1}{8} \times \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix} \quad (17)$$

The gradient matrix \mathbf{G}_n can be calculated as:

$$\mathbf{G}_n = \begin{bmatrix} \mathbf{D}_x \mathbf{n} & \mathbf{D}_y \mathbf{n} \end{bmatrix} \quad (18)$$

where the matrices \mathbf{D}_x and \mathbf{D}_y are derived from filters such as the above. Because the noise is zero-mean, the expected

value of \mathbf{G}_n is:

$$E(\mathbf{G}_n) = \begin{bmatrix} 0 & 0 \end{bmatrix} \quad (19)$$

and the expected gradient covariance matrix becomes:

$$\begin{aligned} E(\mathbf{C}_n) &= E(\mathbf{G}_n^T \mathbf{G}_n) \\ &= E \left(\begin{bmatrix} \mathbf{n}^T \mathbf{D}_x^T \mathbf{D}_x \mathbf{n} & \mathbf{n}^T \mathbf{D}_x^T \mathbf{D}_y \mathbf{n} \\ \mathbf{n}^T \mathbf{D}_y^T \mathbf{D}_x \mathbf{n} & \mathbf{n}^T \mathbf{D}_y^T \mathbf{D}_y \mathbf{n} \end{bmatrix} \right) \end{aligned} \quad (20)$$

where the first entry can be deduced as:

$$\begin{aligned} E(\mathbf{C}_n)_{1,1} &= E(\mathbf{n}^T \mathbf{D}_x^T \mathbf{D}_x \mathbf{n}) \\ &= E(\text{tr}(\mathbf{D}_x \mathbf{n} \mathbf{n}^T \mathbf{D}_x^T)) \\ &= \sigma^2 \text{tr}(\mathbf{D}_x \mathbf{D}_x^T) \end{aligned} \quad (21)$$

and similarly we have:

$$\begin{aligned} E(\mathbf{C}_n)_{1,2} &= \sigma^2 \text{tr}(\mathbf{D}_y \mathbf{D}_x^T), \quad E(\mathbf{C}_n)_{2,1} = \sigma^2 \text{tr}(\mathbf{D}_x \mathbf{D}_y^T) \\ E(\mathbf{C}_n)_{2,2} &= \sigma^2 \text{tr}(\mathbf{D}_y \mathbf{D}_y^T) \end{aligned}$$

The value of the trace $\text{tr}(\mathbf{D}\mathbf{D}^T)$ depends upon the specific filter used in (18). It can be shown that if we choose (16) or (17), the expected \mathbf{C}_n will have the form:

$$E(\mathbf{C}_n) = \begin{bmatrix} \xi N^2 \sigma^2 & 0 \\ 0 & \xi N^2 \sigma^2 \end{bmatrix} \quad (22)$$

where $\xi = \frac{1}{2}$ if we use filters (16), and for filters in (17) $\xi = \frac{3}{16}$.

Now consider how the value of s_1 changes when a clean image \mathbf{g} is corrupted by the noise image denoted by \mathbf{n} . The gradient matrix of the noisy image $\hat{\mathbf{g}}$ would become:

$$\hat{\mathbf{G}} = \mathbf{G} + \mathbf{G}_n \quad (23)$$

Since \mathbf{G} is deterministic, the expected $\hat{\mathbf{C}}$ would have the form:

$$\begin{aligned} E(\hat{\mathbf{C}}) &= E(\hat{\mathbf{G}}^T \hat{\mathbf{G}}) \\ &= \mathbf{G}^T \mathbf{G} + E(\mathbf{G}_n^T \mathbf{G}_n) + 2\mathbf{G}^T E(\mathbf{G}_n) \\ &= \mathbf{V} \begin{bmatrix} s_1^2 + \xi N^2 \sigma^2 & 0 \\ 0 & s_2^2 + \xi N^2 \sigma^2 \end{bmatrix} \mathbf{V}^T \end{aligned} \quad (24)$$

So on average the singular value \hat{s}_1 of the noisy image can approximately be written as:

$$\hat{s}_1 \approx \sqrt{s_1^2 + \xi N^2 \sigma^2} \quad (25)$$

This equation tells us that \hat{s}_1 is determined by both s_1 and σ^2 . Given a fixed σ^2 , the value of \hat{s}_1 drops as s_1 gets decreased, or say when the image \mathbf{g} is more blurry. Unfortunately, \hat{s}_1 is also monotonically increasing with the noise variance σ^2 . So our definition of a sharpness metric must be a modification of \hat{s}_1 , as described below. It is useful to note that if σ^2 is sufficiently large, \hat{s}_1 becomes approximately proportional to the standard deviation σ :

$$\hat{s}_1 \approx \xi^{\frac{1}{2}} N \sigma \quad (26)$$

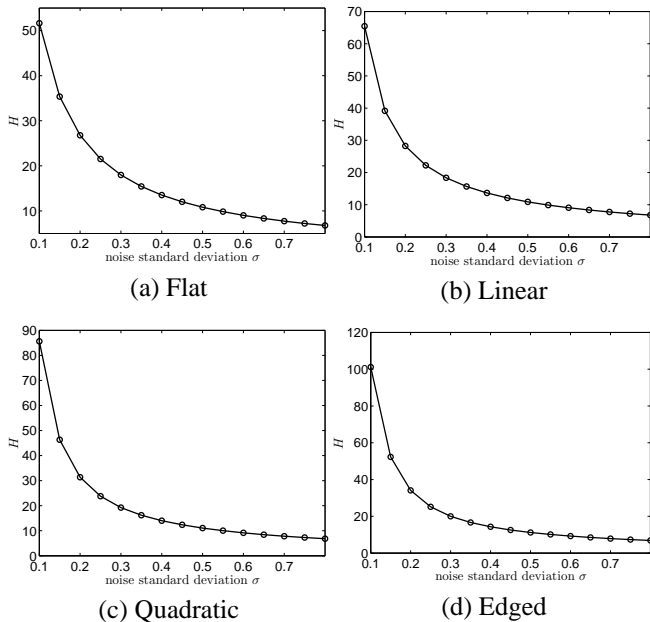


Fig. 3. Plots of the average H versus the noise deviation σ in Monte-Carlo simulations for different patches, where $\epsilon = 5 \times 10^{-4}$ and the pixel intensity range is $[0, 1]$. 100 different noise realizations were used for each σ to get the average H . The patch size $N = 7$.

3. SHARPNESS METRIC H

The above analysis demonstrates that \hat{s}_1 itself cannot be used directly as a sharpness metric in the presence of both noise and blur. To alleviate this problem, we define the sharpness metric H as follows:

$$H = \frac{\hat{s}_1}{\epsilon + \sigma^2} \quad (27)$$

where ϵ is a fixed, small positive constant¹. Here we assume that the noise variance σ^2 is known, or at least can be estimated. For a fixed σ^2 , the behavior of H is basically the same as \hat{s}_1 . As σ^2 rises, however, the value of H drops to zero, as desired.

Monte-Carlo simulations are carried out to examine the behavior of H in different image patches. We add a realization of WGN with σ ranging from 0.1 to 0.8 to the test patch, and calculate the metric H . For each σ , this process is repeated with 100 independent noise realizations to get the average sharpness metric. ϵ is set to be 5×10^{-4} , and the true standard deviation σ is used in calculating H . The plots are illustrated in Fig. 3, which shows that the sharpness metric H is decreasing as the variance of noise rises. In Section 4 we present experimental validation of H using more general natural images with noise, blur, and compression artifacts.

¹We note that $\hat{s}_1 = s_1$ and $H = \frac{s_1}{\epsilon}$ if no noise is present.




Fig. 4. The probability density function of sharpness metric H for white Gaussian noise. (a) Density functions with different window size N , where $\sigma = 1/\sqrt{2}$, $\epsilon = 0$ (b) Density functions with different standard deviation σ , where $N = 9$ and $\epsilon = 0$.

Before describing further experimental results, we pause to note that when noise is present, H is of course a random variable. If the variance of the noise is assumed to be given, then a particular measured value of H can be used to decide how much “true content” the image contains as compared to noise. Said another way, H can be thought of as a rough indicator of the signal to noise ratio. In order for this concept to be useful, it is necessary to know the actual statistical distribution of H . Unfortunately, this distribution is in general quite complicated. But happily, it can be derived in closed form when the image is pure WGN (or approximately when the noise is overwhelming the signal). This is still useful as a baseline value for the statistic H . More specifically, based on results from [9] we can derive this density as:

$$f_H(h) = \frac{1}{2(N^2 - 2)!} e^{-\frac{h^2}{2\mu^2}} \frac{h^{2\beta-1}}{\mu^{2\beta}} \left[2 \left(\frac{h}{\mu} \right)^{2\beta} e^{-\frac{h^2}{2\mu^2}} - 2^\beta \left(2\beta - \frac{h^2}{\mu^2} \right) \gamma \left(\beta, \frac{h^2}{2\mu^2} \right) \right] \quad (28)$$

where

$$\mu = \frac{\epsilon^{\frac{1}{2}}}{\sigma + \epsilon/\sigma}, \quad \beta = \frac{N^2 - 1}{2} \quad (29)$$

$$\gamma(\alpha, x) = \int_0^x t^{\alpha-1} e^{-t} dt$$

Plots of this density with different window sizes and different noise variances are shown in Fig.4. (Due to lack of space, we do not describe the detailed derivation in this paper.)

4. EXPERIMENTS

In this section we illustrate further experimental results for the proposed no-reference sharpness metric H . The experimental measurement strategy we take is the same for all the experiments. Namely, the image is first divided into M non-overlapping blocks, and then H_i is calculated for each block




Fig. 5. Plots of \bar{H} versus smoothing kernel deviation σ_b . WGN with $\sigma_b = 10, 10.1, 10.2$ was added respectively to the test image *Lena*. In calculating local H , the window size $N = 16$, $\epsilon = 1$ and the noise standard deviation estimated through MAD method were used.

i. Finally we use the average $\bar{H} = \frac{1}{M} \sum_{i=1}^M H_i$ as the metric for the whole image. It is worth mentioning that in practice, when calculating H_i , if the standard deviation of noise is not known, it too can be estimated using a variety of methods. In the experiments here we employ the median absolute deviation (MAD) method [10] to measure the standard deviation σ of Gaussian noise globally, and assume that this value is unchanged through all the blocks.

4.1. Simulations

First, we examine the validity of H in measuring blur for natural images with random noise. The test gray image is *Lena* (512×512). A Gaussian blur kernel (of size 10×10), whose standard deviation σ_b ranges from 0.4 to 2.0, is applied to the image respectively, and WGN with variance σ^2 is added after each blurring process. The plots of \bar{H} are shown in Fig. 5. The block size is $N = 16$, and $\epsilon = 1^2$.

Next, we keep the blur kernel fixed and increase the variance of WGN. Results are shown in Fig. 6. It is observed that both blur and noise decrease the value of the proposed sharpness H , as desired.

4.2. JPEG2000 experiments

As an extension of the application of our metric H , here we use it to compute the quality of JPEG2000 compressed images from the LIVE database [11]. The specific images used are shown in Fig. 7. H is calculated from the luminance channel of the images. All the parameters for H are the same as in Section 4.1 above, and we assume no random noise is involved. The recently introduced and popularized image quality metric SSIM [12] is also calculated for comparison.

From the results shown in Fig. 8, it can be seen that the rate distortion curves using H (or say \bar{H} across the whole image) and SSIM as the distortion metric respectively behave

²The value of ϵ is higher compared with the experiments before, because the pixel intensity range here is [0, 255].




Fig. 6. Plots of \bar{H} versus Gaussian deviation noise standard deviation σ . The test image is *Lena*, and the smoothing kernel $\sigma_b = 1.2$. In calculating local H , the window size $N = 16$ and $\epsilon = 1$.



Fig. 7. Images from the LIVE database for JPEG2000 test. Their size is 512×768 .

in a similar fashion. This may indicate that the no-reference metric H is also able to capture the visual quality of the images in the presence of compression artifacts. In addition, it is observed that the metric H has a larger range of values, compared with SSIM (which basically computes the fidelity, and is always between 0 and 1). This is because H indicates the signal to noise ratio as we described before. Images with more detail will get higher value in H . By looking at the estimated \bar{H} of different test images at the same bit rate, we can see that the value for the *Buildings* image, which is more detailed, is higher than *Lighthouse*; and *Monarch*, whose background is blurry, is always the smallest.

5. CONCLUSION

In this paper, we proposed a no-reference sharpness metric H based on image gradients. Experiments show that it can quantify the amount of both blur and random noise, and thus, can capture the change of visual quality in many image processing applications.

As described in Section 3, the metric H actually serves as an indicator of the signal to noise ratio. However, it requires a prior estimation of noise variance. Further study is already under way to devise an image content metric based on an extension of H , which does not need any prior knowledge, and can be utilized in solving the parameter optimization problem for image restoration algorithms, such as denoising [13].




Fig. 8. Plots of mean SSIM and \bar{H} versus bit rate (bpp) for the JPEG2000 compressed images *Buildings*, *Monarch* and *Lighthouse*. In calculating local H , $N = 16$, $\epsilon = 1$ and $\sigma = 0$.

6. REFERENCES

- [1] C. F. Batten, "Autofocusing and astigmatism correction in the scanning electron microscope," in *M.Phil. Thesis, University of Cambridge*, August 2000.
- [2] J. Caviedes and F. Oberti, "A new sharpness metric based on local kurtosis, edge and energy information," in *Signal Processing: Image Communication*, 2004, vol. 19, pp. 147–161.
- [3] X. Marichal, W. Y. Ma, and H. J. Zhang, "Blur determination in the compressed domain using DCT information," in *Proceedings of the International Conference on Image Processing*, Kobe, Japan, October 1999, vol. 2, pp. 386–390.
- [4] P. Marziliano, F. Dufaux, S. Winkler, and T. Ebrahimi, "A no-reference perceptual blur metric," in *Proceedings of the International Conference on Image Processing*, Rochester, NY, 2002, vol. 3, pp. 57–60.
- [5] S. Varadarajan and L. J. Karam, "An improved perception-based no-reference objective image sharpness metric using iterative edge refinement," in *Proceedings of the 15th IEEE International Conference on Image Processing*, October 2008, pp. 401–404.
- [6] R. Ferzli and L. J. Karam, "A no reference objective sharpness metric using Riemannian tensor," in *Third International Workshop on Video Processing and Quality Metrics for Consumer Electronics*, Scottsdale, Arizona, January 2007, pp. 25–26.
- [7] R. Ferzli and L. J. Karam, "No-reference objective wavelet-based noise immune image sharpness metric," in *IEEE International Conference on Image Processing*, September 2005, pp. 405–408.
- [8] X. Feng and P. Milanfar, "Multiscale principal components analysis for image local orientation estimation," in *Proceedings of the 36th Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, CA, November 2002, vol. 1, pp. 478–482.
- [9] A. Edelman, "Eigenvalues and condition numbers of random matrices," *SIAM Journal on Matrix Analysis and Applications*, vol. 9, no. 4, pp. 543–560, October 1988.
- [10] F. R. Hampel, "The influence curve and its role in robust estimation," *Journal of the American Statistical Association*, vol. 69, pp. 383–393, 1974.
- [11] H. R. Sheikh, Z. Wang, L. Cormack, and A. C. Bovik, "LIVE image quality assessment database," <http://live.ece.utexas.edu/research/quality>.
- [12] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, April 2004.
- [13] X. Zhu and P. Milanfar, "A no-reference metric of true image content and its application in denoising," *In preparation*.