

Gray and Color Image Contrast Enhancement by the Curvelet Transform

Jean-Luc Starck, Fionn Murtagh, Emmanuel J. Candès, and David L. Donoho

Abstract—We present in this paper a new method for contrast enhancement based on the curvelet transform. The curvelet transform represents edges better than wavelets, and is therefore well-suited for multiscale edge enhancement. We compare this approach with enhancement based on the wavelet transform, and the Multiscale Retinex. In a range of examples, we use edge detection and segmentation, among other processing applications, to provide for quantitative comparative evaluation. Our findings are that curvelet based enhancement out-performs other enhancement methods on noisy images, but on noiseless or near noiseless images curvelet based enhancement is not remarkably better than wavelet based enhancement.

Index Terms—Contrast enhancement, curvelets, ridgelets, wavelets.

I. INTRODUCTION

BECAUSE some features are hardly detectable by eye in an image, we often transform images before display. Histogram equalization is one the most well-known methods for contrast enhancement. Such an approach is generally useful for images with poor intensity distribution. Since edges play a fundamental role in image understanding, one good way to enhance the contrast is to enhance the edges. For example, we can add to the original image its Laplacian ΔI ($I' = I + \gamma \Delta I$, where I' is the enhanced image and γ is a parameter). Only features at the finest scale are enhanced (linearly). For a high γ value, only the high frequencies are visible. Multiscale edge enhancement [15] can be seen as a generalization of this approach, taking all resolution levels into account.

In color images, objects can exhibit variations in color saturation with little or no correspondence in luminance variation. Several methods have been proposed in the past for color image enhancement [14]. The retinex concept was introduced by Land [7] as a model for human color constancy. The single scale retinex (SSR) method [6] consists of applying the following transform to each band i of the color image:

$$R_i(x, y) = \log(I_i(x, y)) - \log(F(x, y) * I_i(x, y)) \quad (1)$$

Manuscript received April 26, 2002; revised February 6, 2003. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Eric L. Miller.

J.-L. Starck is with the CEA-Saclay, DAPNIA/SEDI-SAP, Service d'Astrophysique, F-91191 Gif sur Yvette, France (e-mail: jstarck@cea.fr).

F. Murtagh is with the School of Computer Science, Queen's University Belfast, Belfast BT7 1NN, Ireland (e-mail: f.murtagh@qub.ac.uk).

E. J. Candès is with the Department of Applied Mathematics, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: emmanuel@acm.caltech.edu).

D. L. Donoho is with the Department of Statistics, Stanford University, Stanford, CA 94305 USA (e-mail: donoho@stat.stanford.edu).

Digital Object Identifier 10.1109/TIP.2003.813140

where $R_i(x, y)$ is the retinex output, $I_i(x, y)$ is the image distribution in the i th spectral band, F is a Gaussian function, and $*$ is convolution. A gain/offset is applied to the retinex output which clips the highest and lowest signal excursions. This can be done by k -sigma clipping. The retinex method is efficient for dynamic range compression, but does not provide good tonal rendition [10]. The Multiscale Retinex (MSR) combines several SSR outputs to produce a single output image which has both good dynamic range compression and color constancy (color constancy may be defined as the independence of the perceived color from the color of the light source [8], [9]), and good tonal rendition [5]. The MSR can be defined by

$$R_{MSR_i} = \sum_{j=1}^N w_j R_{i,j} \quad (2)$$

with

$$R_{i,j}(x, y) = \log(I_i(x, y)) - \log(F_j(x, y) * I_i(x, y)). \quad (3)$$

N is the number of scales, $R_{i,j}$ is the i th spectral component of the MSR output, and w_j is the weight associated with the scale j . The Gaussian F_j is given by

$$F_j(x, y) = K \exp\left(-\frac{r^2}{c_j^2}\right) \quad (4)$$

where c_j defines the width of the Gaussian. In [5], three scales were recommended with c_j values equal respectively to 15, 80, 250, and all weights w_j fixed to $1/N$. These parameters may however be image dependent, and automatic parameter estimation by a genetic algorithm was proposed in [9].

The Multiscale Retinex introduces the concept of multiresolution for contrast enhancement. It performs dynamic range compression and can be used for different image processing goals. Improvements of the algorithm have been presented in [1], leading to better color fidelity.

MSR softens the strongest edges and keeps the faint edges almost untouched. The opposite approach was proposed by Velde [15] in using the wavelet transform for enhancing the faintest edges and keeping untouched the strongest. The strategies are different, but both methods allow the user to see details which were hardly distinguishable in the original image, by reducing the ratio of strong features to faint features.

The wavelet approach [15] consists of first transforming the image using the dyadic wavelet transform (two directions per scale). The gradient $G_{j,k}$ at scale j and at pixel location k is calculated at each scale j from the wavelet coefficients $w_{j,k}^{(h)}$ and $w_{j,k}^{(v)}$ relative to the horizontal and vertical wavelet bands:




Fig. 1. Enhanced coefficients versus original coefficients. Parameters are $m = 30$, $c = 3$, and $p = 0.5$.

$G_{j,k} = \sqrt{(w_{j,k}^{(h)})^2 + (w_{j,k}^{(v)})^2}$. Then the two wavelet coefficients at scale j and at pixel position k [i.e., $k = (x_1, x_2)$] are multiplied by $y(G_{j,k})$, where y is defined by

$$\begin{aligned} y(x) &= \left(\frac{m}{c}\right)^p & \text{if } |x| < c \\ y(x) &= \left(\frac{m}{|x|}\right)^p & \text{if } c \leq |x| < m \\ y(x) &= 1 & \text{if } |x| \geq m. \end{aligned} \quad (5)$$

Three parameters are needed: p , m and c . p determines the degree of nonlinearity in the nonlinear rescaling of the luminance, and must be in $[0, 1]$. Coefficients larger than m are not modified by the algorithm. The c parameter corresponds to the noise level. Fig. 1 shows the modified wavelet coefficients versus the original wavelet coefficients for a given set of parameters ($m = 30$, $c = 3$ and $p = 0.5$). Finally, the enhanced image is obtained by the inverse wavelet transform from the modified wavelet coefficients. For color images, a similar method can be used, but by calculating the overall multiscale gradient $\Gamma_{j,k}$ from the multiscale gradient of the three L , u , v components: $\Gamma_j(i) = \sqrt{\|G_{j,k}^L\|^2 + \|G_{j,k}^u\|^2 + \|G_{j,k}^v\|^2}$. All wavelet coefficients at scale j and at position k are multiplied by $y(\Gamma_{j,k})$, the enhanced \tilde{L} , \tilde{u} , \tilde{v} components are reconstructed from the modified wavelet coefficients, and the $(\tilde{L}, \tilde{u}, \tilde{v})$ image is transformed into an RGB image. More details can be found in [15].

Wavelet bases present some limitations, because they are not well adapted to the detection of highly anisotropic elements, such as alignments in an image, or sheets in a cube. Recently, other multiscale systems have been developed, which include in particular ridgelets [2] and curvelets [3], [12], and these are very different from wavelet-like systems. Curvelets and ridgelets take the form of basis elements which exhibit very high directional sensitivity and are highly anisotropic. The curvelet transform uses the ridgelet transform in its digital implementation. We first describe the ridgelet and the curvelet transforms, and then we show how contrast enhancement can be obtained from the curvelet coefficients. Following that, we present a number of evaluations of the use of wavelet- and curvelet-based enhancement.

II. CONTRAST ENHANCEMENT USING THE CURVELET TRANSFORM

A. Ridgelet Transform

The two-dimensional continuous ridgelet transform in \mathbf{R}^2 can be defined as follows [2]. We pick a smooth univariate func-

tion $\psi: \mathbf{R} \rightarrow \mathbf{R}$ with sufficient decay and satisfying the admissibility condition

$$\int |\hat{\psi}(\xi)|^2 / |\xi|^2 d\xi < \infty \quad (6)$$

where \hat{x} denotes the Fourier transform of x . Equation (6) holds if, say, ψ has a vanishing mean $\int \psi(t) dt = 0$. We will suppose a special normalization about ψ so that $\int_0^\infty |\hat{\psi}(\xi)|^2 \xi^{-2} d\xi = 1$.

For each $a > 0$, each $b \in \mathbf{R}$ and each $\theta \in [0, 2\pi)$, we define the bivariate ridgelet $\psi_{a,b,\theta}: \mathbf{R}^2 \rightarrow \mathbf{R}$ by

$$\psi_{a,b,\theta}(x_1, x_2) = a^{-1/2} \cdot \psi((x_1 \cos \theta + x_2 \sin \theta - b)/a). \quad (7)$$

A ridgelet is constant along lines $x_1 \cos \theta + x_2 \sin \theta = \text{const.}$ Transverse to these ridges it is a wavelet.

Fig. 2 graphs a few ridgelets with different parameter values. The top right, bottom left and right panels are obtained after simple geometric manipulations of the upper left ridgelet, namely rotation, rescaling, and shifting.

Given an integrable bivariate function $f(x)$, we define its ridgelet coefficients by

$$\mathcal{R}_f(a, b, \theta) = \int \bar{\psi}_{a,b,\theta}(x) f(x) dx$$

where \bar{x} denotes the conjugate of x . We have the exact reconstruction formula

$$f(x) = \int_0^{2\pi} \int_{-\infty}^{\infty} \int_0^{\infty} \mathcal{R}_f(a, b, \theta) \psi_{a,b,\theta}(x) \frac{da}{a^3} db \frac{d\theta}{4\pi} \quad (8)$$

valid a.e. for functions which are both integrable and square integrable.

Ridgelet analysis may be construed as wavelet analysis in the Radon domain. Recall that the Radon transform of an object f is the collection of line integrals indexed by $(\theta, t) \in [0, 2\pi) \times \mathbf{R}$ given by

$$Rf(\theta, t) = \int f(x_1, x_2) \delta(x_1 \cos \theta + x_2 \sin \theta - t) dx_1 dx_2 \quad (9)$$

where δ is the Dirac function. Then the ridgelet transform is precisely the application of a 1-D wavelet transform to the slices of the Radon transform where the angular variable θ is constant and t is varying.

This viewpoint strongly suggests developing approximate Radon transforms for digital data. This subject has received considerable attention over the past decades since the Radon transform naturally appears as a fundamental tool in many fields of scientific investigation. Our implementation follows a widely used approach in the literature of medical imaging and is based on fast Fourier transforms. The key component is to obtain approximate digital samples from the Fourier transform on a polar grid, i.e., along lines going through the origin in the frequency plane. Fig. 3 (left) represents the flowgraph of the ridgelet transform. We will not detail this approach further here, and instead refer the reader to [12].

The ridgelet transform of a digital array of size $n \times n$ is an array of size $2n \times 2n$ and hence introduces a redundancy factor equal to 4.

Local Ridgelet Transforms

Speaking in engineering terms, one might say that the ridgelet transform is well-adapted for picking linear structures of about

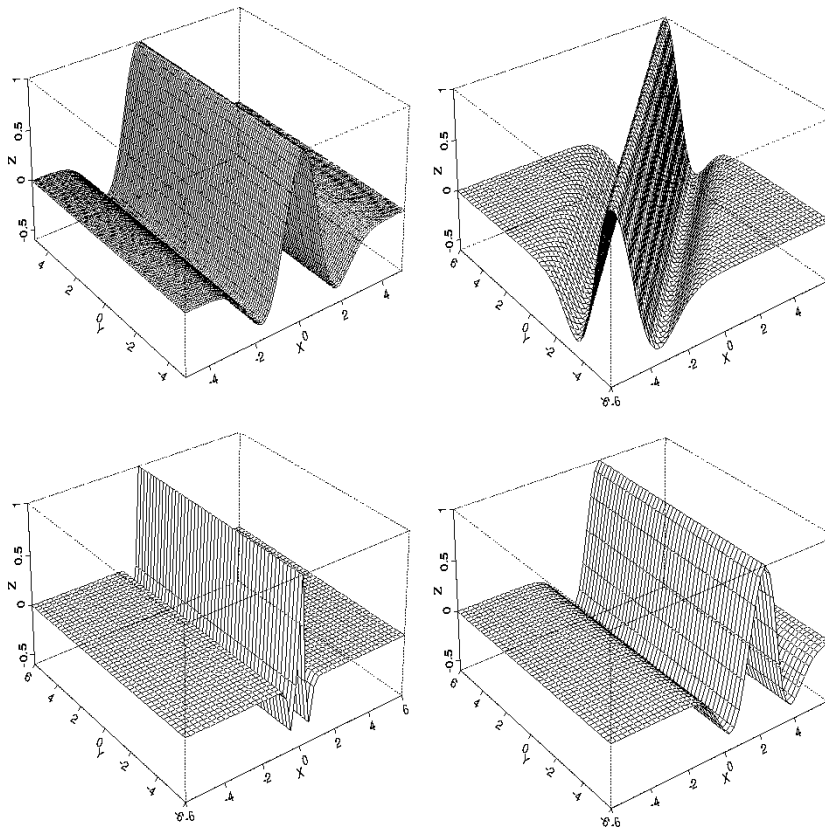


Fig. 2. A few ridgelets.

the size of the image. However, interesting linear structures, e.g., line segments, may occur at a wide range of scales. Following a well-established tradition in time-frequency analysis, the opportunity arises of developing a pyramid of ridgelet transforms. We may indeed apply classical ideas such as recursive dyadic partitioning, and thereby construct dictionaries of windowed ridgelets, renormalized and transported to a wide range of scales and locations.

To make things more explicit we consider the situation at a fixed scale. The image is decomposed into smoothly overlapping blocks of side length b pixels in such a way that the overlap between two vertically adjacent blocks is a rectangular array of size b by $b/2$; we use overlap to avoid blocking artifacts. For an n by n image, we count $2n/b$ such blocks in each direction. The partitioning introduces redundancy, since a pixel belongs to 4 neighboring blocks. More details on a possible implementation of the digital ridgelet transform can be found in [12]. Taking the ridgelet transform of these smoothly localized data is what we call the local ridgelet transform.

B. Curvelet Transform

The idea of curvelets [3] is to represent a curve as a superposition of functions of various lengths and widths obeying the scaling law $width \approx length^2$. This can be done by first decomposing the image into subbands, i.e., separating the object into a series of disjoint scales. Each scale is then analyzed by means of a local ridgelet transform.

Curvelets are based on multiscale ridgelets combined with a spatial bandpass filtering operation to isolate different scales.

This spatial bandpass filter nearly kills all multiscale ridgelets which are not in the frequency range of the filter. In other words, a curvelet is a multiscale ridgelet which lives in a prescribed frequency band. The bandpass is set so that the curvelet length and width at fine scales are related by a scaling law $width \approx length^2$ and so the anisotropy increases with decreasing scale like a power law. There is a very special relationship between the depth of the multiscale pyramid and the index of the dyadic subbands; the side length of the localizing windows is doubled at every other dyadic subband, hence maintaining the fundamental property of the curvelet transform which says that elements of length about $2^{-j/2}$ serve for the analysis and synthesis of the j th subband $[2^j, 2^{j+1}]$. While multiscale ridgelets have arbitrary dyadic length and arbitrary dyadic widths, curvelets have a scaling obeying $width \approx length^2$. Loosely speaking, the curvelet dictionary is a subset of the multiscale ridgelet dictionary, but which allows reconstruction.

In our opinion the “à trous” subband filtering algorithm is especially well-adapted to the needs of the digital curvelet transform. The algorithm decomposes an n by n image I as a superposition of the form

$$I(x, y) = c_J(x, y) + \sum_{j=1}^J w_j(x, y)$$

where c_J is a coarse or smooth version of the original image I and w_j represents “the details of I ” at scale 2^{-j} . See [13] for more information. Thus, the algorithm outputs $J + 1$ subband arrays of size $n \times n$. (The indexing is such that, here, $j = 1$ corresponds to the finest scale, i.e., high frequencies.)

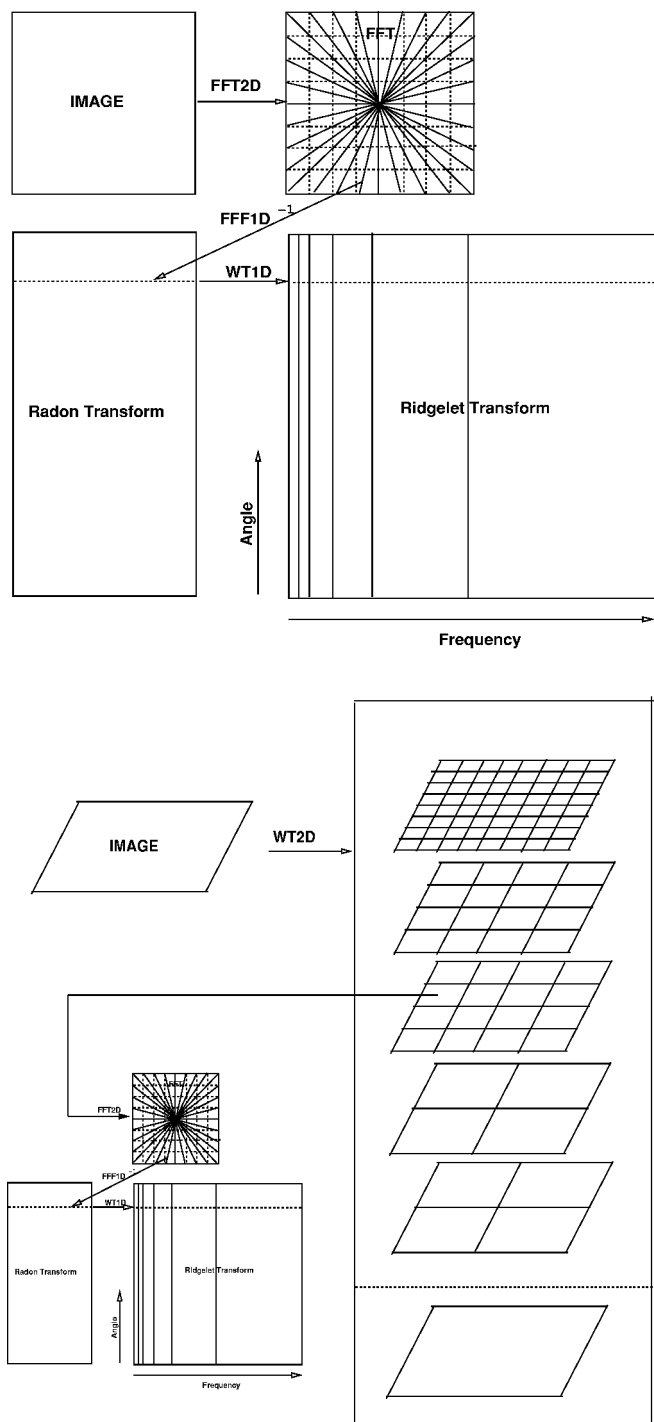


Fig. 3. Top, ridgelet transform flowgraph. Each of the $2n$ radial lines in the Fourier domain is processed separately. The 1-D inverse FFT is calculated along each radial line followed by a 1-D nonorthogonal wavelet transform. In practice, the one-dimensional wavelet coefficients are directly calculated in the Fourier space. Bottom, curvelet transform flowgraph. The figure illustrates the decomposition of the original image into subbands followed by the spatial partitioning of each subband (i.e., each subband is decomposed into blocks). The ridgelet transform is then applied to each block.

As a side comment, we note that the coarse description of the image c_J is not processed. We used the default value $b_{\min} = 16$ pixels in our implementation. Fig. 3 (right) gives an overview of the organization of the algorithm.

This implementation of the curvelet transform is redundant. The redundancy factor is equal to $16J + 1$ whenever J scales are employed. Finally, the method enjoys exact reconstruction and stability, because each step of the transform is both invertible and stable.

III. CONTRAST ENHANCEMENT USING THE CURVELET TRANSFORM

Since the curvelet transform is well-adapted to represent images containing edges, it is a good candidate for edge enhancement. Curvelet coefficients can be modified in order to enhance edges in an image. A function y_c must be defined which modifies the values of the curvelet coefficients. It could be a function similar to the one defined for the wavelet coefficients [15] [see (5)]. This function however gives rise to the drawback amplifying the noise (linearly) as well as the signal of interest. We introduce explicitly the noise standard deviation σ in the equation

$$\begin{aligned}
 y_c(x, \sigma) &= 1 && \text{if } x < c\sigma \\
 y_c(x, \sigma) &= \frac{x - c\sigma}{c\sigma} \left(\frac{m}{c\sigma}\right)^p + \frac{2c\sigma - x}{c\sigma} && \text{if } x < 2c\sigma \\
 y_c(x, \sigma) &= \left(\frac{m}{x}\right)^p && \text{if } 2c\sigma \leq x < m \\
 y_c(x, \sigma) &= \left(\frac{m}{x}\right)^s && \text{if } x \geq m.
 \end{aligned} \tag{10}$$

Here, p determines the degree of nonlinearity and s introduces dynamic range compression. Using a nonzero s will enhance the faintest edges and soften the strongest edges at the same time. c becomes a normalization parameter, and a c value larger than 3 guaranties that the noise will not be amplified. The m parameter is the value under which coefficients are amplified. This value depends obviously on the pixel values inside the curvelet scale. Therefore, we found it necessary to derive the m value from the data. Two options are possible:

- m can be derived from the noise standard deviation ($m = K_m \sigma$) using an additional parameter K_m . The advantage is that K_m is now independent of the curvelet coefficient values, and therefore much easier for a user to set. For instance, using $c = 3$ and $K_m = 10$ amplifies all coefficients with a SNR between 3 and 10.
- m can also be derived from the maximum curvelet coefficient M_c of the relative band ($m = lM_c$, with $l < 1$). In this case, choosing for instance $c = 3$ and $l = 0.5$, we amplify all coefficients with an absolute value between 3σ and half the maximum absolute value of the band.

The first choice allows the user to define the coefficients to be amplified as a function of their signal to noise ratio, while the second one gives an easy and general way to fix the m parameter independently of the range of the pixel values. Fig. 4 shows the curve representing the enhanced coefficients versus the original coefficients for two sets of parameters.

The curvelet enhancement method for grayscale images consists of the following steps.

- 1) Estimate the noise standard deviation σ in the input image I .




Fig. 4. Enhanced coefficients versus original coefficients. Left, parameters are $m = 30$, $c = 3$, $s = 0$, and $p = 0.5$. Right, parameters are $m = 30$, $c = 3$, $s = 0.6$, and $p = 0.5$.

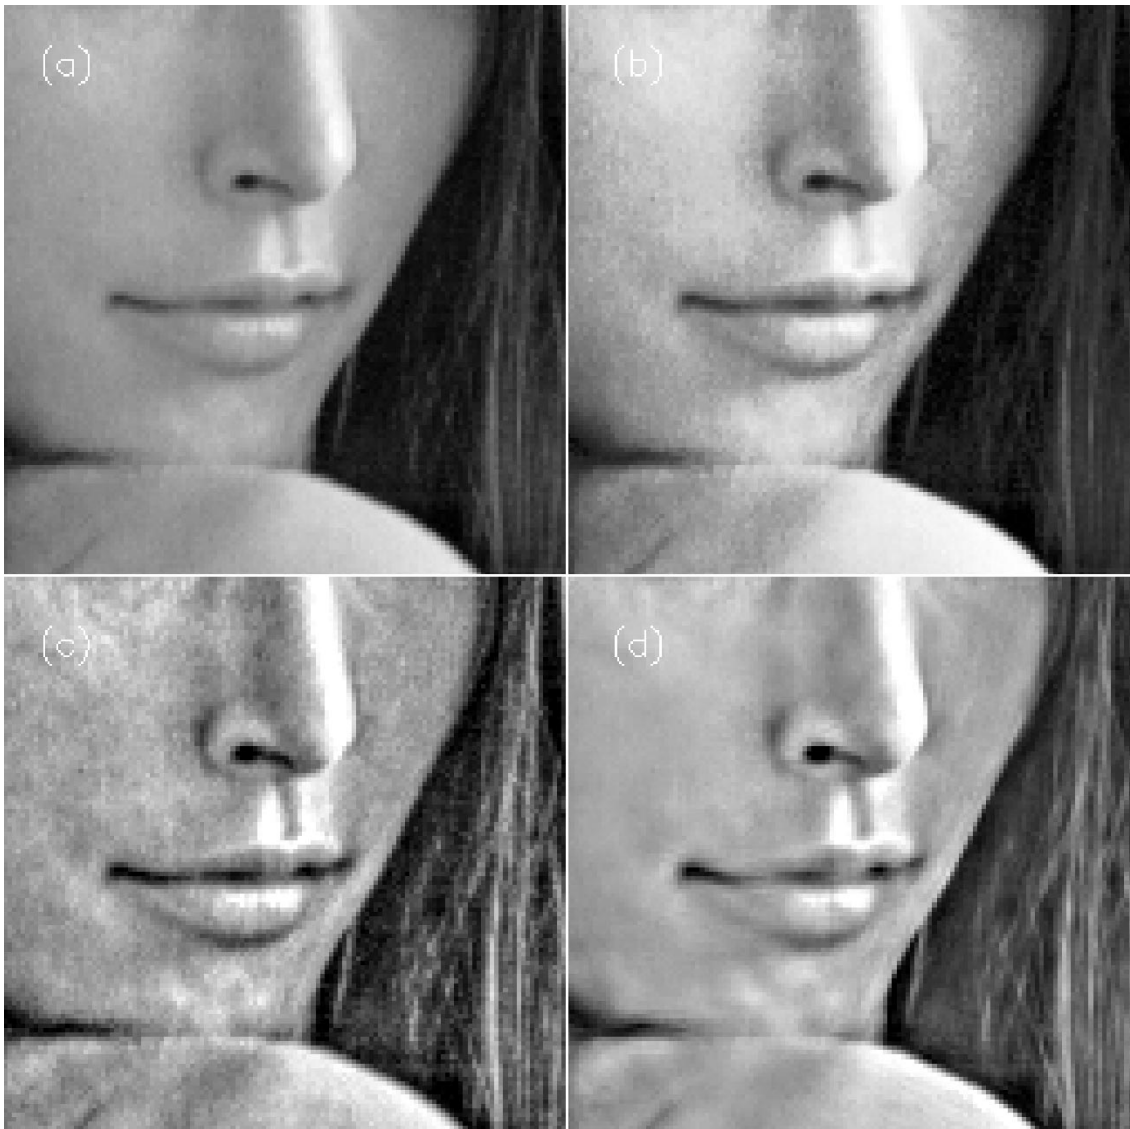


Fig. 5. Top: part of Lena image and its histogram equalization. Bottom: enhanced image by the wavelet transform and the curvelet transform.

2) Calculate the curvelet transform of the input image. We get a set of bands w_j , each band w_j contains N_j coefficients and corresponds to a given resolution level.

3) Calculate the noise standard deviation σ_j for each band j of the curvelet transform (see [12] for more details on this step).

4) For each band j do

- Calculate the maximum M_j of the band.
- Multiply each curvelet coefficient $w_{j,k}$ by $y_c(|w_{j,k}|, \sigma_j)$.

5) Reconstruct the enhanced image from the modified curvelet coefficients.




Fig. 6. Top, grayscale image, and bottom, curvelet enhanced image.

For color images, we apply first the curvelet transform to the three components L , u , v . For each curvelet coefficient, we calculate $e = \sqrt{c_L^2 + c_u^2 + c_v^2}$, where (c_L, c_u, c_v) are, respectively, the curvelet coefficients of the three components, and the modified coefficients are obtained by: $(\tilde{c}_L, \tilde{c}_u, \tilde{c}_v) = (y_c(e, \sigma)c_L, y_c(e, \sigma)c_u, y_c(e, \sigma)c_v)$.

Values in the enhanced components can be larger than the authorized upper limit (in general 255), and we found it necessary

to add a final step to our method, which is a gain/offset selection applied uniformly to the three color subimages, as described in [6].

Examples

Fig. 5 shows the results of, respectively, histogram equalization, wavelet and curvelet enhancement, using the standard Lena test image. No noise was added to the image used, implying




Fig. 7. Top, color image (Kodak picture of the day 14/05/02) and retinex method. Bottom, multiscale retinex method and curvelet edge enhancement.

small levels only of quantization noise present. The better result seen here for the curvelet enhancement (Fig. 5 bottom right) is in part due to the Velde method [15] used in the wavelet-based method over-enhancing small noise levels.

Fig. 6 shows the results for the enhancement of a grayscale satellite image (parameters were $c = 3$, $l = p = 0.5$ and $s = 0$).

Fig. 7 shows the results for the enhancement of a color image (Kodak image of the day 14/05/01) by the retinex (same parameters), the multiscale retinex and the curvelet multiscale edge enhancement methods. Fig. 8 shows the results for the enhancement of a color image (Kodak image of the day 11/12/01). These examples present some evidence for the benefits of curvelet




Fig. 8. Top, color image (Kodak picture of the day 11/12/01), and bottom, curvelet enhanced image.

enhancement. Small, aligned features are preserved well. Note however that better color fidelity can be obtained for the MSR image by using the color restoration algorithm described in [1].

In summary, the results of these three figures indicate that the curvelet based enhancement approach works well. In the next section, we will evaluate it relative to other enhancement approaches, and in particular wavelet based enhancement.

IV. EVALUATION

A. Evaluation Methodology

Image enhancement quality is difficult to assess. Considerable literature exists relative to image quality estimation [11], [4]. However, this is most often in the context of image compression where the problem is to estimate the distortion or the

loss of information, with criteria other than PSNR (peak signal to noise ratio), because PSNR does not reflect errors in the way that the human vision system does. For image enhancement, the goal is to introduce distortion, in such a way that some low level or low contrast features can easily be seen by a human operator. A subjective assessment approach is simply to present images enhanced by different methods, as we did in the previous section, and to let a domain expert judge the best result.

In order to have an object quality criterion, we will make the following assumption: between two edge enhancement techniques, the better one will be that which produces the best results for standard vision processing tasks, such as segmentation or edge detection. We do not claim that image enhancement should be applied before carrying out a segmentation or an edge detection (other pre-processing steps such as filtering




Fig. 9. Left, image containing a number of bars, and right, bar edge image.

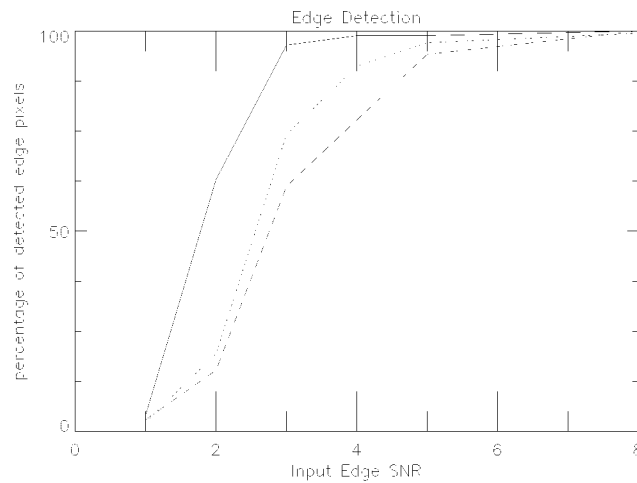


Fig. 10. Percentage of detected edge pixels versus the edge SNR using a Canny edge detector on the wavelet enhanced image (dashed line), the curvelet enhanced image using Velde's function enhancement (dotted line), and the curvelet enhanced image using the new function enhancement (continuous line).

are certainly more appropriate), but we consider that if an image enhancement method improves the human performance for analyzing a scene, it should do the same for a machine-based vision approach. We describe two experiments in the following, providing some measure of objectivity for comparison of results, using edge detection and segmentation. Finally, we return to the issue of the limits of curvelet versus wavelet enhancement.

B. Edge Detection

Fig. 9 consists of an artificial image containing a number of bars. The intensity is constant along each individual bar; from left to right, the intensities of the six vertical bars (these are in fact thin rectangles which are 20 pixels wide and 150 pixels long, having a 30° angle with the x -axis) are respectively equal to 1, 2, 3, 4, 5, 8. The noise standard deviation is 1.

We ran the wavelet and the curvelet methods on this simulated image. The curvelet method was applied twice, once with Velde's enhancement function and once with the proposed enhancement function. Then we applied a Canny edge detector on the three enhanced images. We estimated the noise in the three edge images from pixels outside the bars, and considered as edges all pixels with a value larger than five times the noise standard deviation. Knowing the right edges (they were extracted by applying the Canny edge detector to the original noise free image; see Fig. 9 right), we derived the percentage

of recovered edge pixels: this is 54.77% for the wavelet-based image, 64.66% for the curvelet enhanced image using Velde's function enhancement, and 73.91% for the curvelet enhanced image using the new function enhancement. As each bar has a different intensity level, we can also derive the percentage of recovered edge pixels as a function of the edge signal to noise ratio (SNR). Fig. 10 shows such a curve. This gives the percentage of detected edge pixels versus the edge SNR using a Canny edge detector on i) the wavelet enhanced image (dashed line), ii) the curvelet enhanced image using Velde's function enhancement, and iii) the curvelet enhanced image using the new function enhancement (continuous line).

These results are clearly in favor of the curvelet transform.

C. Segmentation

Contrast enhancement can facilitate user interpretation of an image, or it can help in automated interpretation. Here, we will use segmentation as an important processing goal. We will use a grayscale 512×512 Lena test image on account of its smooth and edge regions.

The alternative contrast enhancement approaches used are: i) histogram equalization, using the algorithm in the IDL image processing package, ii) wavelet coefficient enhancement, as described in Section I above, and iii) curvelet transform based enhancement, as described in Section II-B above.




Fig. 11. Marginal density histograms (binsize = 3) of original Lena image (top left), histogram equalized image (top right), wavelet enhanced (bottom left), and curvelet enhanced (bottom right).



Fig. 12. A five-segment result, using a Markov Potts model, of the original image.



Fig. 13. A five-segment result, using a Markov Potts model, of the histogram equalized image.

Fig. 11 shows the marginal densities of these images. Histogram equalization essentially destroys information relative to pixel classification through marginal density fitting. With histogram equalization, image quantization remains feasible, of course, but it is clear from Fig. 12 that possibly useful information is lost. Wavelet enhancement (bottom left panel in Fig. 12) also smooths out information. Only the curvelet enhancement (bottom right panel in Fig. 12) retains marginal density fidelity to the original image marginal density (upper left panel).

To investigate the quality of segmentations carried out on these images, we used a five-component Gaussian fit, based on a Markov random field model with neighborhood 3×3 , and with a Potts/Ising spatial model. The spatial influence parameter, ϕ , did not differ greatly among these results. We found, for the

original and histogram-equalized images, and the wavelet- and curvelet-enhanced images, respective values of: 0.72, 0.72, 0.63, and 0.73. We also determined, as measures of model fit, pseudo-likelihood information criterion values, with limited explanatory capability in this instance.

The segmentation results are shown in Figs. 12–15. In the histogram equalized result (Fig. 13) edge information is destroyed: cf. details of the big cap feather. The wavelet-enhanced result (Fig. 14) does very well in edge regions: cf. details of the cap feather, and hair. However some injustice is done to the smooth regions. The curvelet enhancement (Fig. 15) performs well in edge regions (feather, background) while simultaneously respecting smooth areas. Overall, from the points of view of marginal density, and also spatial segmentation, we find the



Fig. 14. A five-segment result, using a Markov Potts model, of the wavelet-enhanced image.



Fig. 15. A five-segment result, using a Markov Potts model, of the curvelet-enhanced image.

curvelet transform enhancement method to provide a better result which is simultaneously "close" to the original input image.

V. CONCLUSION

A number of properties, respected by the curvelet filtering described here, are important for contrast stretching.

1) Reconstruct the enhanced image from the modified curvelet coefficients. Noise must not be amplified in enhancing edges.

2) Reconstruct the enhanced image from the modified curvelet coefficients. Colors should not be unduly modified. In the multiscale retinex, for example, a tendency toward increased grayness is seen. This is not the case using curvelets.

However color restoration could also be carried out in a final step, as proposed for the multiscale retinex [1]. This should improve the final image quality.

3) Reconstruct the enhanced image from the modified curvelet coefficients. It is very advantageous if block effects do not occur. Block overlapping is usually not necessary in curvelet-based contrast enhancement, unlike in the case of noise filtering.

A range of further examples can be seen at <http://www-stat.stanford.edu/~jstarck/contrast.html>.

Our conclusions are as follows.

1) The curvelet and wavelet enhancement functions take account very well of image noise.

2) As evidenced by the experiments with the curvelet transform, there is better detection of noisy contours than with other methods.

3) For noise-free images, there is not a great deal to be gained by curvelet enhancement over wavelet enhancement since the enhancement function tends toward Velde's approach in such weak noise cases. Contours and edges are detected quite adequately by wavelets in such situations.

ACKNOWLEDGMENT

The authors would like to thank the referees for some very helpful comments on the original version of the manuscript.

REFERENCES

- [1] K. Barnard and B. Funt, "Investigations into multi-scale retinex," in *Color Imaging: Vision and Technology*. New York: Wiley, 1999, pp. 9–17.
- [2] E. J. Candès, "Harmonic analysis of neural networks," *Appl. Comput. Harmon. Anal.*, vol. 6, pp. 197–218, 1999.
- [3] E. J. Candès and D. L. Donoho, "Curvelets—A surprisingly effective nonadaptive representation for objects with edges," in *Curve and Surface Fitting: Saint-Malo 1999*, A. Cohen, C. Rabut, and L. L. Schumaker, Eds. Nashville, TN: Vanderbilt University Press, 1999.
- [4] I. Hontsch and L. J. Karam, "Adaptive image coding with perceptual distortion control," *IEEE Trans. Image Processing*, vol. 11, pp. 213–222, Mar. 2002.
- [5] D. J. Jobson, Z. Rahman, and G. A. Woodell, "A multi-scale retinex for bridging the gap between color images and the human observation of scenes," *IEEE Trans. Image Processing*, vol. 6, pp. 965–976, July 1997.
- [6] D. J. Jobson, Z. Rahman, and G. A. Woodell, "Properties and performance of a center/surround retinex," *IEEE Trans. Image Processing*, vol. 6, pp. 451–462, Mar. 1997.
- [7] E. Land, "Recent advances in retinex theory," *Vis. Res.*, vol. 26, no. 1, pp. 7–21, 1986.
- [8] H. A. Mallot, *Computational Vision: Information Processing in Perception and Visual Behavior*. Cambridge, MA: MIT Press, 2000.
- [9] C. Munteanu and A. Rosa, "Color image enhancement using evolutionary principles and the retinex theory of color constancy," in *Proc. 2001 IEEE Signal Processing Society Workshop on Neural Networks for Signal Processing XI*, 2001, pp. 393–402.
- [10] Z. Rahman, D. J. Jobson, and G. A. Woodell, "Multi-scale retinex for color image enhancement," in *IEEE Int. Conf. Image Processing*, 1996.
- [11] D. Schilling and P. C. Cosman, "Image quality evaluation based on recognition times for fast image browsing applications," *IEEE Trans. Multimedia*, vol. 4, no. 3, pp. 320–331, 2002.
- [12] J. L. Starck, E. Candès, and D. L. Donoho, "The curvelet transform for image denoising," *IEEE Trans. Image Processing*, vol. 11, pp. 131–141, June 2002.
- [13] J. L. Starck, F. Murtagh, and A. Bijaoui, *Image Processing and Data Analysis: The Multiscale Approach*. Cambridge, U.K.: Cambridge Univ. Press, 1998.
- [14] A. Toet, "Multiscale color image enhancement," in *Proc. SPIE Int. Conf. Image Processing and Its Applications*, 1992, pp. 583–585.
- [15] K. V. Velde, "Multi-scale color image enhancement," in *Proc. Int. Conf. Image Processing*, vol. 3, 1999, pp. 584–587.



Jean-Luc Starck received the Ph.D. degree from the University Nice-Sophia Antipolis and the Habilitation degree from University Paris XI.

He was a Visitor at the European Southern Observatory (ESO) in 1993 and in the Statistics Department, Stanford University, Stanford, CA, in 2000. He has been a Researcher at CEA since 1994. His research interests include image processing, multiscale methods, and statistical methods in astrophysics. He is the author of two books entitled *Image Processing and Data Analysis: The Multiscale Approach* (Cam-

bridge, U.K.: Cambridge University Press, 1998) and *Astronomical Image and Data Analysis* (Berlin, Germany: Springer, 2002).



Fionn Murtagh received the B.A. and B.A.I. degrees in mathematics and engineering science, and the M.Sc. degree in computer science, all from Trinity College, Dublin, Ireland. He received the Ph.D. degree in mathematical statistics from the Université P. & M. Curie, Paris VI, France, and an Habilitation degree from Université L. Pasteur, Strasbourg, France.

Previous posts have included Senior Scientist with the Space Science Department of the European Space Agency, and visiting appointments with the European

Commission's Joint Research Centre, and the Department of Statistics, University of Washington. He is Professor of Computer Science at Queen's University, Belfast, Ireland. He is Editor-in-Chief of *The Computer Journal*.

Dr. Murtagh is a Fellow of the British Computer Society.



Emmanuel J. Candès graduated from the Ecole Polytechnique, France, and received the M.Sc. degree in applied mathematics from the University of Paris VI, France. He received the Ph.D. degree in statistics at Stanford University, Stanford, CA, where David L. Donoho served as his adviser.

He is Assistant Professor of Applied and Computational Mathematics at the California Institute of Technology (Caltech), Pasadena. Prior to joining Caltech, he was an Assistant Professor in statistics at Stanford University, Stanford, CA. His research interests are in

the areas of computational harmonic analysis, approximation theory, statistical estimation, and their applications to signal and image processing and scientific computing.

Dr. Candès is an Alfred P. Sloan Research Fellow.



David L. Donoho received the A.B. degree (*summa cum laude*) in statistics from Princeton University, Princeton, NJ, where his senior thesis adviser was John W. Tukey, and the Ph.D. degree in statistics from Harvard University, Cambridge, MA, where his Ph.D. adviser was Peter Huber.

He is Professor of Statistics at Stanford University. He has previously been a Professor at the University of California, Berkeley, and a Visiting Professor at Université de Paris, as well as a Sackler Fellow at Tel Aviv University. His research interests are in harmonic analysis, image representation, and mathematical statistics.

Dr. Donoho is a member of the U.S.A. National Academy of Sciences and a fellow of the American Academy of Arts and Sciences.