Adaptive Image Contrast Enhancement Using Generalizations of Histogram Equalization

J. Alex Stark

***Abstract*—This paper proposes a scheme for adaptive image contrast enhancement based on a generalization of histogram equalization (HE). HE is a useful technique for improving image contrast, but its effect is too severe for many purposes. However, dramatically different results can be obtained with relatively minor modifications.**

A concise description of adaptive HE is set out, and this framework is used in a discussion of past suggestions for variations on HE. A key feature of this formalism is a "cumulation function," which is used to generate a grey level mapping from the local histogram. By choosing alternative forms of cumulation function one can achieve a wide variety of effects. A specific form is proposed. Through the variation of one or two parameters, the resulting process can produce a range of degrees of contrast enhancement, at one extreme leaving the image unchanged, at another yielding full adaptive equalization.

***Index Terms*—Adaptive histogram equalization, contrast enhancement, histogram equalization, image enhancement.**

I. Introduction

CONTRAST enhancement techniques are used widely in image processing. One of the most popular automatic procedures is histogram equalization (HE) [1], [2]. This is less effective when the contrast characteristics vary across the image. Adaptive HE [3]–[6] (AHE) overcomes this drawback by generating the mapping for each pixel from the histogram in a surrounding window. AHE does not allow the degree of contrast enhancement to be regulated. The extent to which the character of the image is changed is undesirable for many applications. (An example of the severity of AHE is given in Fig. 1.) One suggested method [7] for obtaining a range of effects between full HE and leaving an image unchanged involves blurring the local histogram before evaluating the mapping.

The first aim of this paper is to set out a concise mathematical description of AHE. The second aim is to show that the resulting framework can be used to generate a variety of contrast enhancement effects, of which HE is a special case. This is achieved by specifying alternative forms of a function which we call the *cumulation function*. (Blurring the image histogram can be interpreted in such terms.) The third aim is to suggest one form of cumulation function; this is defined in terms of two parameters, each with a simple interpretation. The procedure which we propose is flexible and can be implemented efficiently.

Use of the Fourier series method of HE [8], [9] for implementing these suggestions is given particular attention.

Manuscript received March 18, 1999; revised October 15, 1999. This work was supported by Christ's College, Cambridge, U.K., through a research fellowship and also by the National Institute of Statistical Sciences, Research Triangle Park, NC. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Scott T. Acton.

The author is with the National Institute of Statistical Sciences (NISS), Research Triangle Park, NC USA (e-mail: stark@niss.org) and also with the Signal Processing Group, Department of Engineering, University of Cambridge, Cambridge, U.K. (e-mail: jas@eng.cam.ac.uk).

Publisher Item Identifier S 1057-7149(00)03918-X.

II. Adaptive Histogram Equalization

The AHE process can be understood in different ways. In one perspective the histogram of grey levels (GL's) in a window around each pixel is generated first. The cumulative distribution of GL's, that is the cumulative sum over the histogram, is used to map the input pixel GL's to output GL's. If a pixel has a GL lower than all others in the surrounding window the output is maximally black; if it has the median value in its window the output is 50% grey.

This section proceeds with a concise mathematical description of AHE which can be readily generalized, and then considers the two main types of modification. The relationship between the equations and different (conceptual) perspectives on AHE, such as GL comparison, might not be immediately clear, but generalizations can be expressed far more easily in this framework.

A. Mathematical Description

AHE can be described using few equations. Although the framework that follows has a number of complex details, these are all important for modification and implementation. It is intended to be a summary statement of AHE rather than an extensive exposition.

To equalize an input image x with quantized GL's scaled between $-1/2$ and $1/2$, we first require an estimate $\hat{h}$ of the local histogram. (Some implementations do not actually evaluate any histograms, but can be said to do so implicitly.) We can start by sifting those pixels in the input image with GL g using the Kronecker delta function $\delta(i, j)$, which equals 1 if $i = j$ and 0 otherwise. Spatial convolution with a rectangular kernel f_w can then be used to find the number of such pixels in a window around each point. It is convenient to scale f_w so that it is unit-volume; the estimate histogram then sums to unity at each point. For a square window of width w, with odd-integer value, this can be written

$$\hat{h}(m,n,g) = \delta(g, x(m,n)) \overset{m,n}{\star} f_w(m,n) \tag{1}$$

$$f_w(m,n) = \begin{cases} w^{-2}, & |m| \le (w-1)/2, \quad |n| \le (w-1)/2 \\ 0, & \text{otherwise.} \end{cases} \tag{2}$$

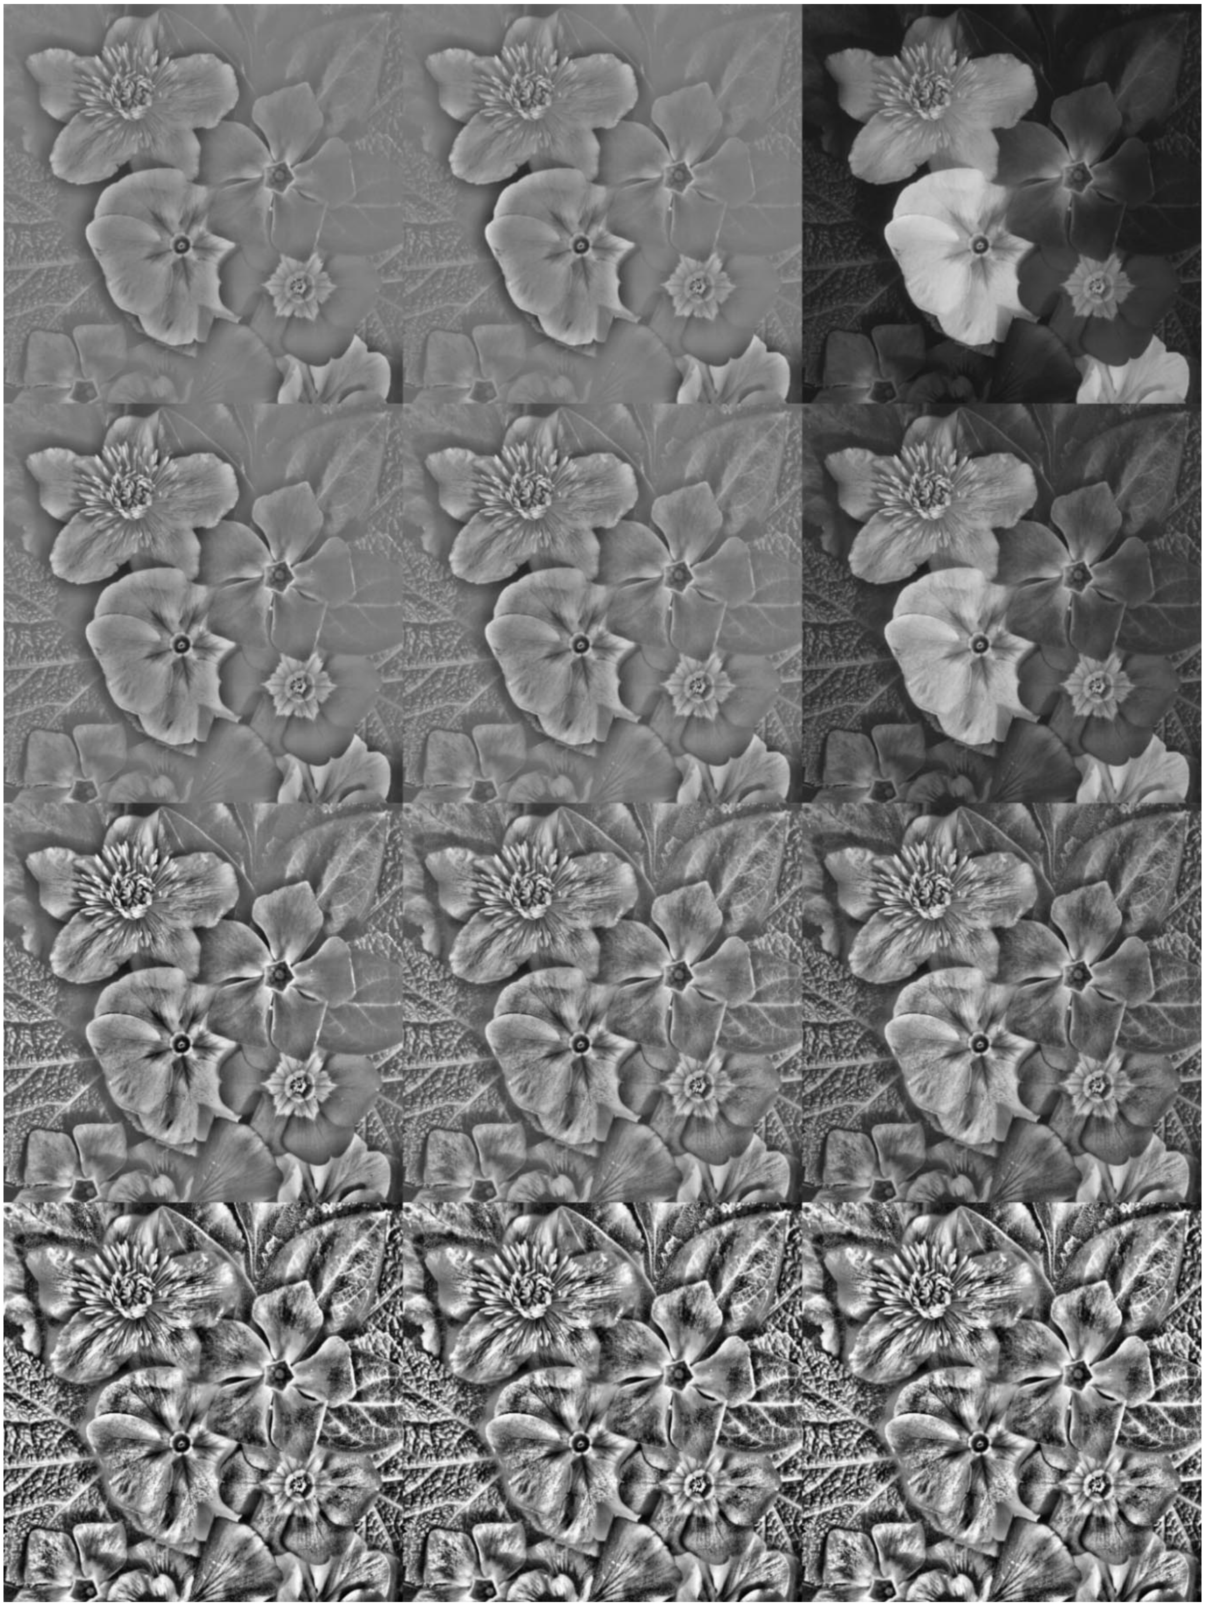

Fig. 1. Three different methods for reducing the effect of histogram equalization (HE). The original spring flowers image (top-right, 512 $\times$ 512 pixels), was equalized using a square window (width 41) to create the image (repeated) at the bottom. HE with Gaussian blurring was used to generate the other images in the first column. From the bottom, the blurring widths were 0 (that is, standard HE), 22, 58, and 90 (for input range 0 to 255). The center-top image was generated by subtracting the local mean with the same window. The signed power-law (SPL) process was used for the other images in the central column; the sequence of α values for the four images was 0, 0.3, 0.65, and 1 from the bottom. These results are very similar to Gaussian blurring: the Gaussian widths were chosen to achieve this. Note the differences in the treatment of the central white flower. Gaussian blurring retains more of the contrast between the darker center and surrounding lighter areas, whereas the power-law process highlights fine detail. The images in the right column were generated using a SPL cumulation function and local-mean replacement. A proportion of the local mean (β) was added back into the result of the power-law process. The replacement proportions were given the same values as the powers $(\beta = \alpha)$.

These equations are explained in more detail in [9]. The output image y is found using

$$y(m,n) = z(m,n,x(m,n)) \tag{3}$$

where z is a spatially varying mapping. In standard HE, the cumulative histogram is used for this. Because $\hat{h}$ sums to unity, z can be constructed using an offset of $1/2$ so that the output ranges from $-1/2$ to $1/2$. We add a third term so that negating the input GL values negates the output:

$$\begin{aligned} z(m,n,g) &= -\frac{1}{2} + \sum_{\gamma<g} \hat{h}(m,n,\gamma) + \frac{1}{2}\hat{h}(m,n,g) \\ &= \frac{1}{2}\sum_{\gamma<g} \hat{h}(m,n,\gamma) - \frac{1}{2}\sum_{\gamma>g} \hat{h}(m,n,\gamma). \end{aligned} \tag{4}$$

This can be generalized using what we call a "cumulation function" f_c

$$z(m,n,g) = \sum_{\gamma} \hat{h}(m,n,\gamma) f_c(g,\gamma). \tag{5}$$

In standard HE, f_c only operates on GL differences

$$f_c(u,v) = f_0(u-v) \tag{6}$$

$$f_0(d) = \begin{cases} 1/2, & B > d > 0 \\ 0, & d = 0 \\ -1/2, & 0 > d > -B. \end{cases} \tag{7}$$

The limit B must be equal to, or greater than, the maximum GL difference. (This limit is exploited in the Fourier series implementation of HE, which employs a periodic cumulation function.) Alternatively, z can be described using convolution over g

$$z(m,n,g) = \hat{h}(m,n,g) \overset{g}{\star} f_0(g). \tag{8}$$

The cumulation function can also be seen as comparing differences in pixel GL's. This approach to AHE has been exploited in rank-based implementations [6], [7].

AHE was applied to a 512-square image (Fig. 1, top-right) with a square window ($w = 41$ pixels) for the weighting function. As is typical with HE, details in the resulting image (Fig. 1, bottom, repeated) are highlighted, but noise is also enhanced. The main drawback with this is that the character of the image has been changed fundamentally. There are many situations in which contrast enhancement is desired but the nature of the image is important.

B. Window Modification

The use of weighting functions f_w for estimating the local histogram (1) other than simple rectangular windows has been repeatedly explored. In one study [6] the difference between conical and rectangular weighting functions was found to be too small to justify the consumption of time. In another article [7] it was argued that the limitation of the dimensions of square windows to odd integers is severe and that the rotational variance was undesirable. Considerable effort was made to develop a tractable algorithm for circular Gaussian weighting functions. A third study considered the task of finding the local histogram as a problem of statistical estimation [9]. It might be possible to develop adaptive filtering processes which are in some sense optimal. Other proposals include varying the window dimensions adaptively across the image [10], or between fields in the Fourier series process [8], and building a neighborhood around each pixel [11]. A recent article applied equalization to connected components which were identified using mathematical morphology [12].

The Fourier series method, further improvements to it, and increases in computer speed may make possible more elaborate filtering and weighting schemes. However, in this research we decided to keep to a square window and concentrate on modifications to the cumulation function. The differences in the results are far more dramatic than those obtained by changing the shape of weighting function. Some comparisons of different contrast effects and window widths are presented in Section IV-A.

C. Gaussian Blurring

One modification of the enhancement process, proposed in [7], is to smooth the local histogram using convolution with a blurring function $b(g)$ such as a Gaussian kernel. This can be described as a modification of (8)

$$z(m,n,g) = \hat{h}(m,n,g) \overset{g}{\star} b(g) \overset{g}{\star} f_0(g). \tag{9}$$

Since the second convolution can be performed first, blurring can also be seen as creating a modified cumulation function. As the width of the Gaussian is increased the convolution of b and f_0 becomes approximately linear over the range of input GL differences. (B must be sufficiently large, as discussed in Section IV-B.) The related case

$$f_c(u,v) = u - v \tag{10}$$

has a simple interpretation. Since the sum of $\hat{h}(m,n,g)$ over g is unity, substituting this into (5) yields

$$z(m,n,g) = g - \sum_{\gamma} \gamma \hat{h}(m,n,\gamma). \tag{11}$$

The second term is the local mean, and so the resulting process subtracts the local mean [9]. The effect of Gaussian blurring is illustrated by the first column of images in Fig. 1. Gaussians of widths (equivalent standard deviations) 22, 58, and 90 GL's were used. The latter width was chosen to match the result of local-mean subtraction (center-top image).

Thus Gaussian blurring, in its strict form, becomes similar to local-mean subtraction as the width of the Gaussian is increased. (The GL's are also scaled because z flattens out.) In [7] the process was modified so that changing the width yielded results between the original and fully equalized images.

III. Modified Cumulation Functions

There are situations in which it is desirable to enhance details in an image without significantly changing its general characteristics. There are also situations in which stronger contrast enhancement is required, but where the effect of standard AHE is too severe. Our aim was to develop a fast and flexible method of adaptive contrast enhancement which could deal with such tasks using few parameters.

The version of Gaussian blurring proposed in [7] is quite effective, but slow to process because a histogram is processed for each pixel. In this study we considered alternative forms of enhancement by directly specifying alternative cumulation functions. The procedure we set out adds flexibility and is simple, and parameter values can be chosen quite easily. A wide variety of contrast enhancement effects can be obtained using the procedure outlined in Section II-A and choosing different cumulation functions. Any reasonable choice can be accommodated by the Fourier series method, and so the challenge is to develop a simple form which produces a desirable range of effects.

A. Signed Power-Law Cumulation

In Section II-C it was shown that local mean subtraction could be obtained using (10). A simple form of cumulation which provides variation between this and standard HE, written in the manner of (6) and (7), is

$$f_c(u,v) = q(u-v,\alpha) \tag{12}$$
$$q(d,\alpha) = \frac{1}{2}\,\mathrm{sign}(d)|2d|^{\alpha} \tag{13}$$

where $0 \le \alpha \le 1$. The special case of $\alpha = 0$ gives HE, and $\alpha = 1$ gives local-mean subtraction. The images in the central column of Fig. 1 show the effect of these values along with $\alpha = 0.65$ and 0.3. This scheme can be used to mitigate the effect of HE, but does not provide a means for slightly increasing the contrast of an image.

B. Local Mean Replacement

More complex forms of cumulation function can be built by superposing components. By substitution into (5) one can see that using

$$f_c(u,v) = u \tag{14}$$

will output the image unchanged. We could find the local mean of the example image by subtracting the top-center image in Fig. 1 from the top-right image. In general we can output the local-mean field by using

$$f_c(u,v) = u - q(u-v,1). \tag{15}$$

Therefore one method for obtaining effects between full HE and retaining the original is to replace a proportion, β, of the local mean

$$f_c(u,v) = q(u-v,\alpha) - \beta q(u-v,1) + \beta u \tag{16}$$

where $0 \le \beta \le 1$. We found empirically that this scheme is effective when α and β are given similar values. The right column of Fig. 1 shows the results for $\beta = \alpha$. In the next section we consider the choice of these parameters in more detail.

C. Selecting α and β

The following theorem is useful in the selection of α and β.

Theorem 1: When applying signed power-law adaptive contrast enhancement with local mean replacement, as set out in (3), (5), (13), and (16), a sufficient condition for ensuring that the output values range between $-1/2$ and $1/2$ is

$$\alpha = \beta. \tag{17}$$

This condition is necessary for practical purposes when spatial convolution (1) is used. (It might be possible to deviate slightly from (17), as explained in the proof.) The condition is sufficient for any local histogram estimation algorithm.

Proof: It is only necessary to consider the case of the upper bound on the output: treatment of the lower bound is the same, mutatis mutandis. First consider the maximum possible output when α and β are given. Since $\alpha \ge 0$, the cumulation function of (13) and (16) is monotonically nondecreasing with u. Combine this with (3) and (5), and suppose that the local histogram is estimated from a reference image x' instead of the input image. In the normal situation $x' = x$ the output may be more tightly bound. The maximum output will arise for a given local histogram when the input $x(m,n) = 1/2$. This is given by

$$y(m,n) = \sum_{\gamma} \hat{h}(m,n,\gamma)\psi(\gamma) \tag{18}$$

where

$$\psi(\gamma) = \frac{1}{2}|1-2\gamma|^{\alpha} + \beta\gamma. \tag{19}$$

The maximum value of ψ for $\gamma < 1/2$ can be seen by differentiating it with respect to γ. Specifically

$$\psi(0) = 1/2 \tag{20}$$
$$\psi'(0) = \beta - \alpha \tag{21}$$
$$\psi''(\gamma) \le 0. \tag{22}$$

Consequently, ψ will have a value greater than $1/2$ if and only if $\beta \ne \alpha$. Equation (17) is therefore a sufficient condition to ensure that $y(m,n) \le 1/2$. If $\beta \ne \alpha$ and the estimate local histogram is concentrated at the corresponding value of γ, then y will exceed $1/2$.

There are two ways in which satisfying (17) exactly might be unnecessary. First, ψ may not be greater than $1/2$ for a valid input GL value. Second, it may not possible to concentrate quite all the histogram at the required GL when $x' = x$. In the case of simple estimation by spatial convolution (1), the pixel at the center of the window contributes to the local histogram through $f_w(0,0)$, but the contribution is almost always small. ■

This theorem means that $\alpha = \beta$ is a safe choice. Out-of-range values can occur in different situations. When local-mean subtraction is applied to an image, a GL value must be more than half the GL range away from the local mean. On the other hand, equalizing an image spreads its values across the full GL range, and so replacing any amount of local mean may push values outside the valid range. In practice contrast enhancement is applied to images with low contrast. Therefore the first situation is encountered less often, and clipping of the output is more likely to occur when α is near 0 and when $\beta > \alpha$.

For a given window width or weighting function, it is possible, by varying α and β, to obtain a wide range of results. This

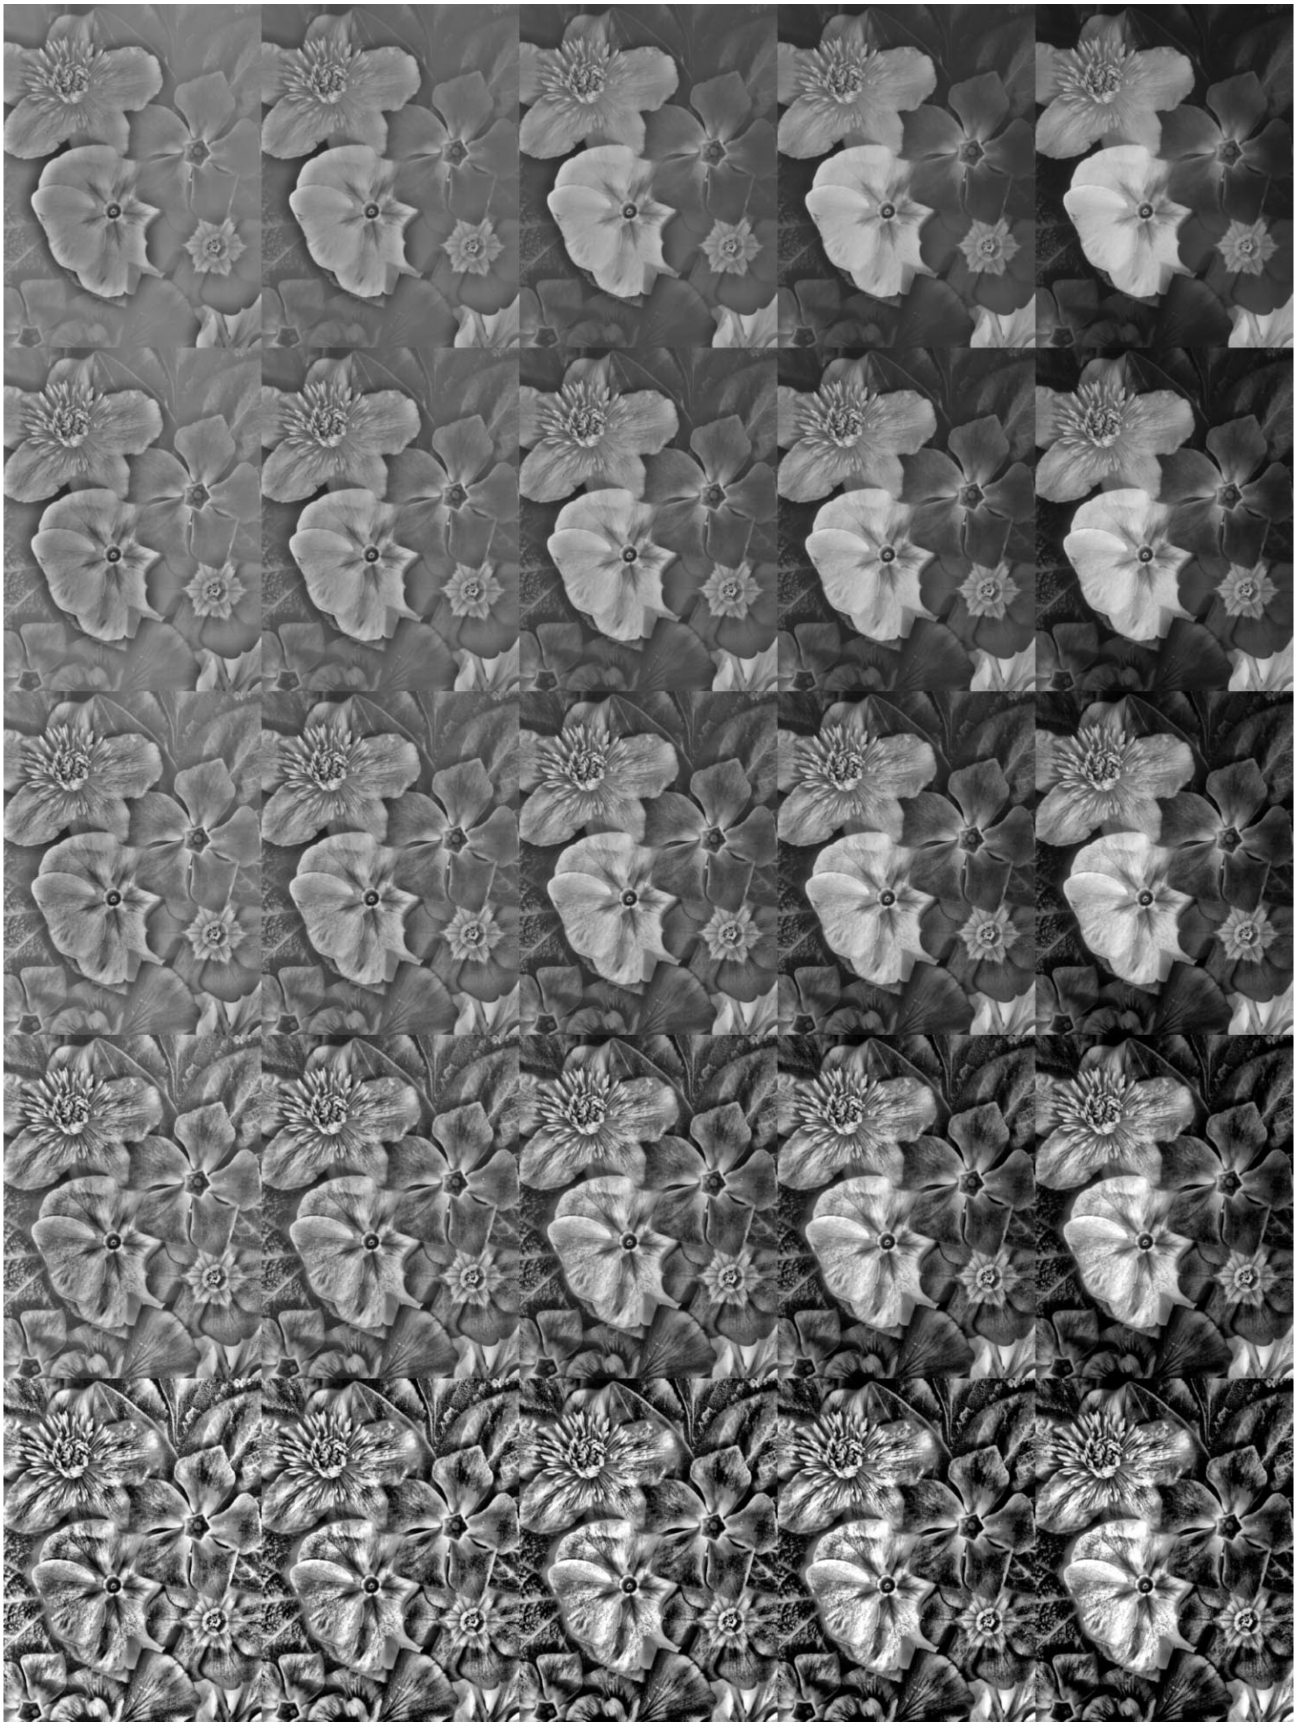

Fig. 2. An array of different enhancement results. Each image was created with a different cumulation function power (α) and proportion of local-mean replacement (β). The top-right image is the original image (cropped), equivalent to using $\alpha = \beta = 1$. The top-left image has had the local mean subtracted ($\alpha = 1, \beta = 0$); the bottom-left image is fully equalized ($\alpha = \beta = 0$). The sequence of α values (one for each row), and β values (one for each column) was: 0, 0.25, 0.5, 0.75, and 1.

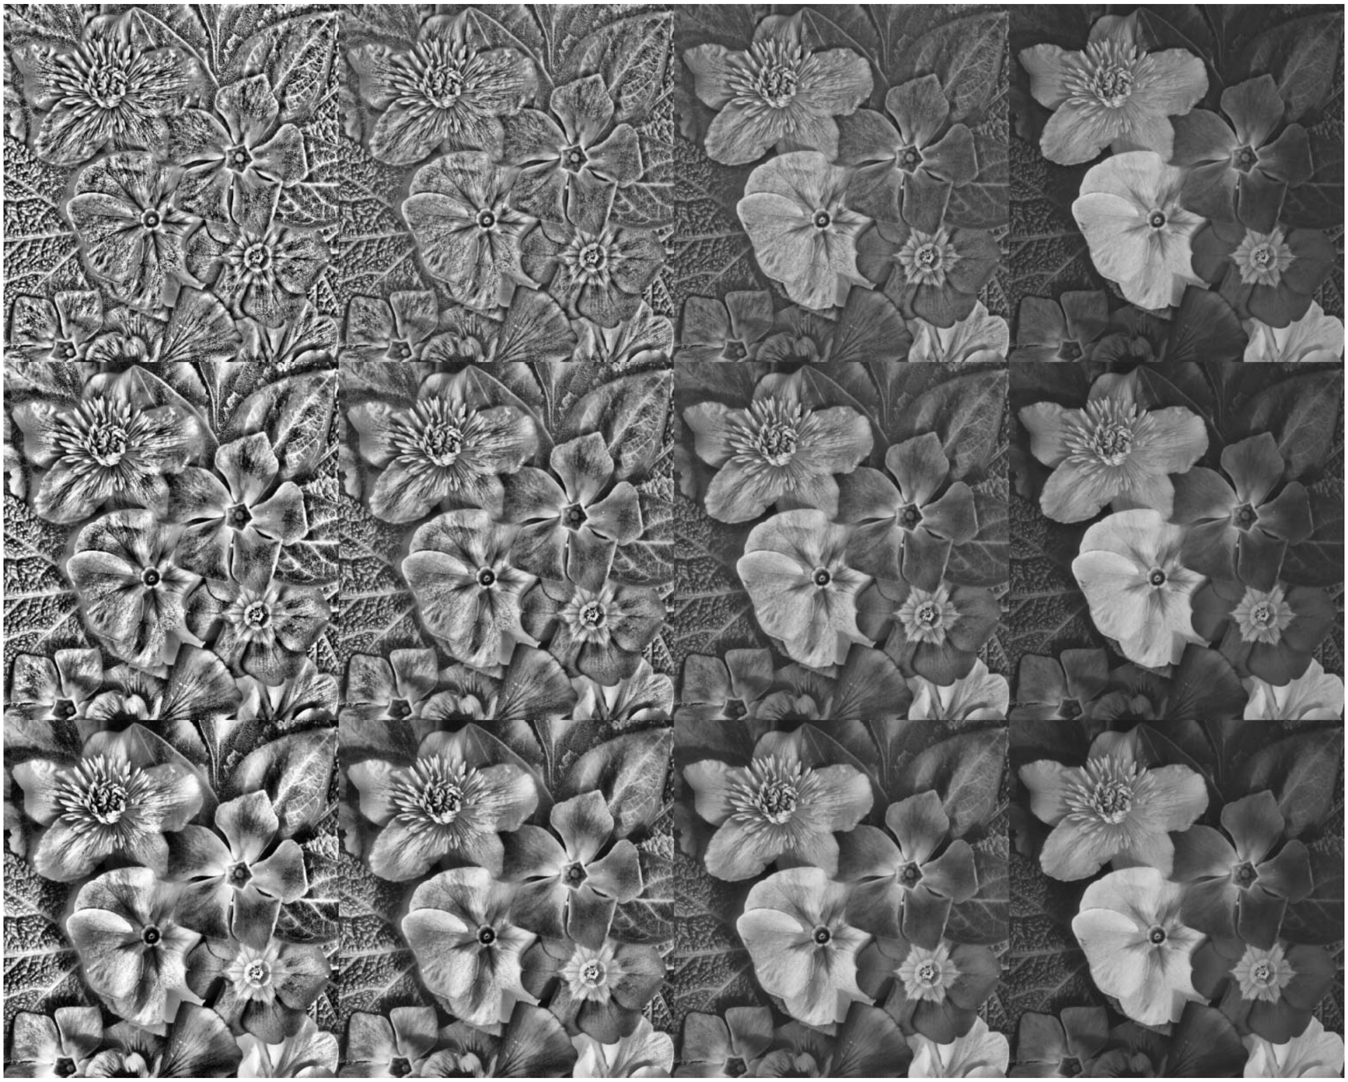

Fig. 3. Combinations of different window widths and contrast effect parameters ($\beta = \alpha$). The images in each row were generated using different α values: 0, 0.15, 0.45, and 0.75 from left to right. For the top row a square window of width 21 was used; the widths for the others were 41 and 81. The effect of changing the width is only small when $\alpha \geq 0.75$ and increases as α is lowered.

is illustrated in Fig. 2. Images on the bottom-left to top-right diagonal show the results for $\alpha = \beta$. The selection of these parameters can be approached in different ways. For example, one might be interested in strong contrast enhancement but wish to soften the effect by setting both parameters near to 0. Alternatively, a small amount of enhancement can be obtained by setting both parameters near to 1. This figure clearly illustrates the effect of choosing different α and β values.

IV. Examples and Methods

A. Further Examples

Although the parameters controlling the spatial estimation process are largely separate from those which determine the GL transformation, it is interesting to look at different combinations. Fig. 3 is an array of images generated with three different window widths and four different contrast effect parameter selections.

Two examples in Fig. 4 show that clipping can be significant. In fact, many of the images in Fig. 2 with $\beta > \alpha$ have been clipped, even when not obvious.

The remaining images in Fig. 4 provide a comparison of results on a test image in the same pattern as Fig. 3. Notice the shadows around edges between regions of different mean level and the ghosting of the bars with the larger window. The bottom half of the images illustrate the results around edges between regions with different contrast. The image generated using a window of width 11 has artifacts. While this choice of window width was deliberately far too small, it is not easy to say what minimum window width is required to avoid such artifacts.

B. Method Details

The results presented here were generated using the Fourier series method. For standard HE, the expansion of f_0 is based on

$$f_0(u-v) = \frac{2}{\pi}\sum_{k=1}^{\infty}\frac{1}{2k-1}\sin\left[(2k-1)\frac{\pi}{2B}(u-v)\right]. \quad (23)$$

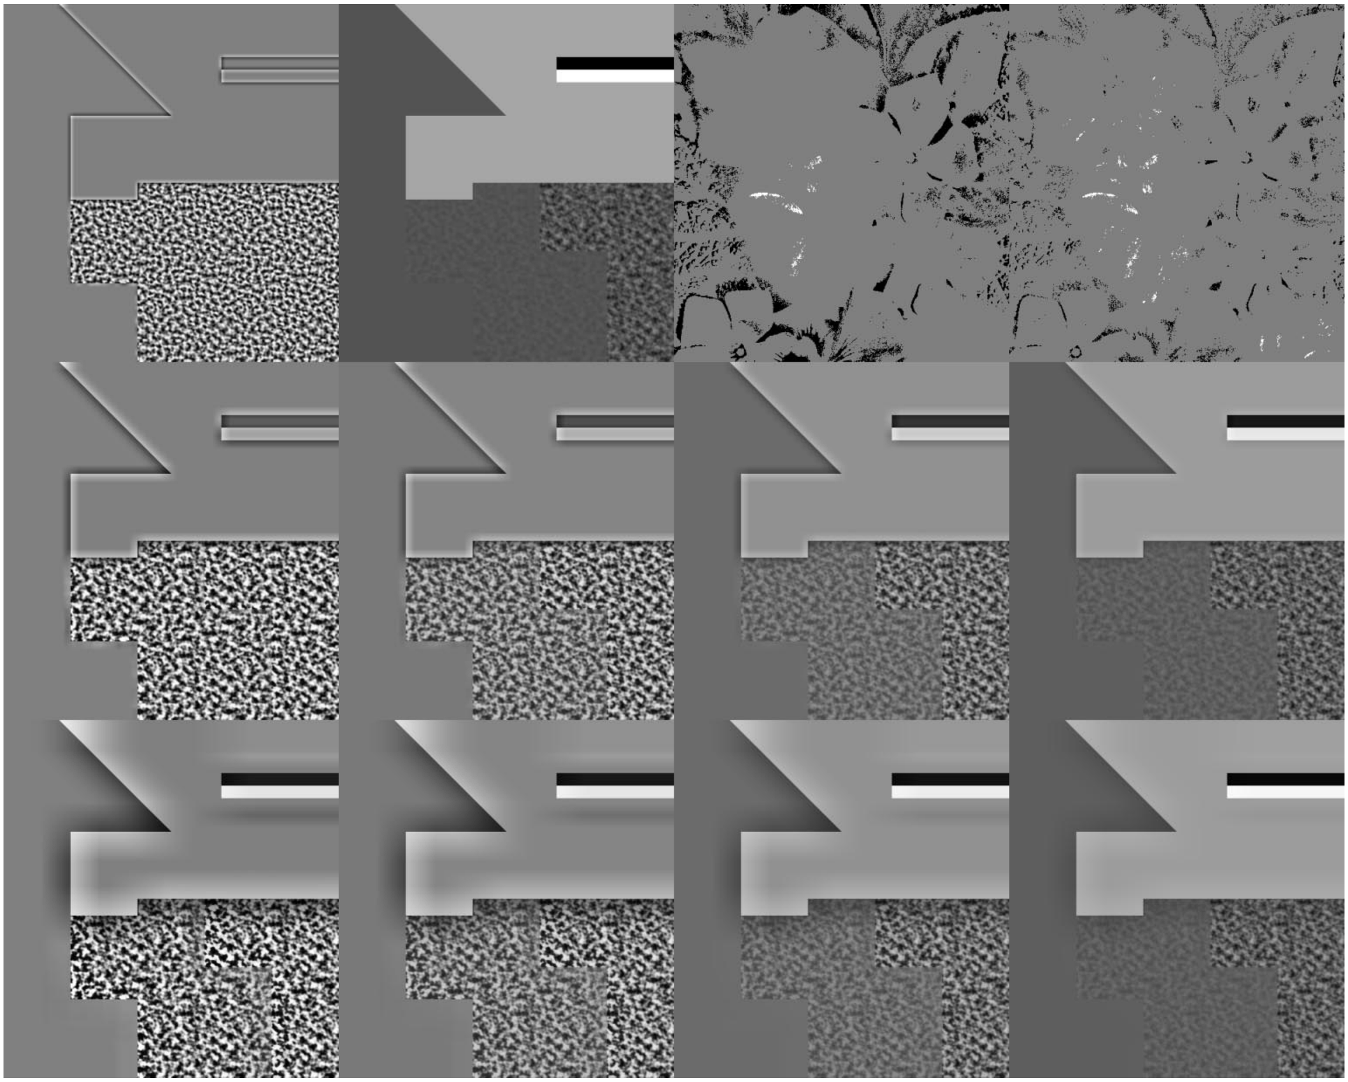

Fig. 4. Test examples. The images in the top right show the clipped pixels (black and white) in two of the images in Fig. 2. The top-right image was generated using $\alpha = 0$ and $\beta = 0.25$; for the other $\alpha = 0.25$ and $\beta = 1$. The second image from the left is a test image. The second and third rows were generated using the same contrast effect parameters as those in Fig. 3, but with window widths 25 and 85. The top-left image was generated using $\alpha = \beta = 0$ and window width 11.

The sine is split into separate terms for u and v; details of the method are contained in [8]. B was set to 1, and 70 terms were used for all examples. This is more terms than would normally be used. (Note that some results such as the the clipping examples in Fig. 4 change slightly with the number of terms.) The transformations of pixels at the edges of the images were generated using the full-sized window with the nearest center.

The alternative cumulation functions were implemented by modifying the Fourier series coefficients. Gaussian blurring was implemented, for example, by multiplying the coefficients by values along a Gaussian, since the Fourier transform of a Gaussian is itself a Gaussian. (The periodic form of the cumulation function in the Fourier series method means that this is an approximation which deteriorates as the Gaussian's width is increased relative to B. It is good for the images in Fig. 1: consider the square-wave of period 4, as given by (23), smoothed by a Gaussian with $\sigma = 90/256$.) The example images for Gaussian blurring were not scaled.

When f_c is a combination of q terms such as in (16), the overall process can be built as the sum of components. It is interesting to note that there are alternative combinations which are exactly equivalent. For example, an alternative to (16) is

$$f_c(u,v) = q(u-v,\alpha) - \beta q(v,1). \tag{24}$$

We used (16), implementing the βu component simply by adding some input to the output. The advantage with this over (24) is that the Fourier series coefficients are smaller when α and β are given values near to 1.

The Fourier series method can be used for any form of f_c, and not just those which are functions of $(u-v)$. Our implementation employs lookup tables for scaled sines and cosines of (scaled) u and v for each term pair. It is convenient to incorporate the series coefficients into the u lookup tables. Thus series coefficients which vary with u can be accommodated. To exploit this one would find the Fourier series over $(u-v)$ for each u. The fact that changes to the cumulation function can be made simply by modifying coefficients means that most of the implementation can be generic: the parts for generating lookup tables, filtering image fields and accumulating the output are basically unchanged.

We consider the Fourier series method to be a "fast" method because it avoids the generation and manipulation of a histogram at each point, and because the computational complexity is proportional to the size of the image and is (largely) independent of the size of contextual region. Other HE methods might be used with modified cumulation functions. In the sampling and interpolation method [6] full histograms are obtained for a subset of pixels. Evaluation of the mappings for each would probably demand the use of a Fourier type procedure.

V. Discussion

The starting point for this paper was a concise mathematical description of AHE. A spatially-varying nonlinear function (3) is used to map input GL's to output GL's. There are two component tasks in generating the mapping. An estimate histogram is found through spatial smoothing (1), and cumulation converts this into the mapping (5). Proposed variations on AHE have generally focussed on one of these two processes. The weighting function f_w has been generalized from the spatially-invariant rectangular form, and modifications have been proposed to the method of converting the estimated histogram into the GL mapping.

We have made two main suggestions. First, the generalized form of adaptive contrast enhancement set out in Section II-A provides considerable flexibility, largely through the cumulation function. Second, simple forms of cumulation functions such as signed power-law with local-mean replacement (16) can yield a wide range of useful effects. In practice, we have found that choosing values of α and β to achieve a desired effect is quite easy. Setting $\beta = \alpha$ is good for many purposes. In cases where the aesthetic quality of the image is important one can start near to $\alpha = \beta = 1$. In cases where strong enhancement is required, but where standard HE is too severe, values around 0.15 are effective. This approach can be useful for making an initial choice which can then be refined. An example is to increase the amount of mean variation by raising the value of β over that of α. Clipping can be a problem, and so this should be done with care.

Forms of cumulation function other than ours may also yield useful results. However, some properties may be important. Specifically, a sigmoidal shape means that larger GL differences $(u - v)$ contribute more to the output. In general $f_c(u, v)$ should probably be nondecreasing with respect to u and nonincreasing with respect to v. Use of (16) means that the process is invariant on image inversion

$$f_c(u, v) = -f_c(-u, -v). \tag{25}$$

The Fourier series method can implement all such forms of adaptive contrast enhancement. Since the changes are mainly to the series coefficients, most of the complexity is in configuration. This method can also process 16-bit images. We have also emphasized the connection with Gaussian blurring. It is a useful technique for dealing with noise. However, for contrast enhancement more generally, we consider our suggested form of cumulation function to be an improvement, both in terms of flexibility and ease of specifying parameters.

The combination of the parameters which control the contrast effect, those controlling the histogram estimation (commonly window size and shape), and those specific to the implementation (such as the number of series terms and cumulation function period) provide a very wide range of effects.

Open-source software implementations of these methods with command-line and graphical interfaces are available online: http://www2.eng.cam.ac.uk/~jas. A plugin for the Gimp program is available online: http://registry.gimp.org.

Acknowledgment

The author wishes to thank those who have contributed to the development of the software, especially K. Turner.

References

[1] A. K. Jain, *Fundamentals of Digital Image Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1989.

[2] W. K. Pratt, *Digital Image Processing*. New York: Wiley, 1978.

[3] D. J. Ketcham, R. Lowe, and W. Weber, "Seminar on image processing," in *Real-Time Enhancement Techniques*, 1976, pp. 1–6. Hughes Aircraft.

[4] R. Hummel, "Image enhancement by histogram transformation," *Comp. Graph. Image Process.*, vol. 6, pp. 184–195, 1977.

[5] V. T. Tom and G. J. Wolfe, "Adaptive histogram equalization and its applications," *SPIE Applicat. Dig. Image Process. IV*, vol. 359, pp. 204–209, 1982.

[6] S. M. Pizer, E. P. Amburn, J. D. Austin, R. Cromartie, A. Geselowitz, T. Greer, B. H. Romeny, J. B. Zimmerman, and K. Zuiderveld, "Adaptive histogram equalization and its variations," *Comp. Vis. Graph. Image Process.*, vol. 39, no. 3, pp. 355–368, 1987.

[7] J. M. Gauch, "Investigations of image contrast space defined by variations on histogram equalization," *CVGIP: Graph. Models Image Process.*, vol. 54, pp. 269–280, July 1992.

[8] J. A. Stark and W. J. Fitzgerald, "An alternative algorithm for adaptive histogram equalization," *Graph. Models Image Process.*, vol. 58, pp. 180–185, Mar. 1996.

[9] ——, "Model-based adaptive histogram equalization," *Signal Process.*, vol. 39, pp. 193–200, 1994.

[10] R. Dale-Jones and T. Tjahjadi, "A study and modification of the local histogram equalization algorithm," *Pattern Recognit.*, vol. 26, no. 9, pp. 1373–1381, 1993.

[11] R. B. Paranjape, W. M. Morrow, and R. M. Rangayyan, "Adaptive-neighborhood histogram equalization for image enhancement," *CVGIP: Graph. Models Image Process.*, vol. 54, pp. 259–267, May 1992.

[12] V. Caselles, J.-L. Lisani, J.-M. Morel, and G. Sapiro, "Shape preserving local histogram modification," *IEEE Trans. Image Processing*, vol. 8, pp. 220–230, Feb. 1999.

J. Alex Stark graduated from the University of Cambridge, Cambridge, U.K., in 1992. He received the Ph.D. degree in statistical signal modeling from the Department of Engineering, University of Cambridge, in 1996.

He has worked on a variety of projects including phylogenetic inference for HIV, ion channel signal analysis, and Bayesian model selection. He recently returned from the National Institute of Statistical Sciences, Research Triangle Park, NC, after taking leave from a research fellowship at Christ's College, Cambridge. Some of his current research interests are: statistical modeling, digital signal processing, image processing, bioinformatics, and computational methods.