

Adaptive hybrid mean and median filtering of high-ISO long-exposure sensor noise for digital photography

Tamer Rabie

UAE University

College of Information Technology

P.O. Box 15551

Al-Ain, United Arab Emirates

E-mail: tamer@cs.toronto.edu

Abstract. *This paper presents a new methodology for the reduction of sensor noise from images acquired using digital cameras at high-International Organization for Standardization (ISO) and long-exposure settings. The problem lies in the fact that the algorithm must deal with hardware-related noise that affects certain color channels more than others and is thus nonuniform over all color channels. A new adaptive center-weighted hybrid mean and median filter is formulated and used within a novel optimal-size windowing framework to reduce the effects of two types of sensor noise, namely blue-channel noise and JPEG blocking artifacts, common in high-ISO digital camera images. A third type of digital camera noise that affects long-exposure images and causes a type of sensor noise commonly known as “stuck-pixel” noise is dealt with by pre-processing the image with a new stuck-pixel prefilter formulation. Experimental results are presented with an analysis of the performance of the various filters in comparison with other standard noise reduction filters. © 2004 SPIE and IS&T. [DOI: 10.1117/1.1668279]*

1 Introduction

With the advent of the inexpensive charge-coupled device (CCD) on a chip (Fig. 1), the wide-spread move from traditional 35 mm film photography to digital photography is becoming increasingly apparent especially with journalists and professional photographers. This has prompted digital camera manufacturers to try to implement most of the legacy techniques common among traditional film cameras such as high-International Organization for Standardization (ISO) film, long exposures, high-speed shutters, etc., into digital cameras. One technique that is of utmost importance to a large community of photographers is the digital camera equivalent to the traditional high-speed silver-based film sensitivity, commonly known as the ISO sensitivity number.

An ISO number that appears on regular camera film packages specifies the speed, or sensitivity, of this type of silver-based film. The higher the number the “faster” or more sensitive the film is to light. Typical ISO speeds for

silver-based film include 100, 200, or 400. Each doubling of the ISO number indicates a doubling in film speed so each of these films is twice as fast as the next fastest. Image sensors used in digital cameras are also rated using equivalent ISO numbers. Just as with film, an image sensor with a lower ISO needs more light for a good exposure than one with a higher ISO. In poorly lit conditions, a longer exposure of the image sensor is needed for more light to enter. This, however, will lead to the acquired images being blurred, unless the scene being imaged is completely still. It is, therefore, better to set the image sensor to a higher ISO setting because this will enhance freezing of scene motion and shooting in low light. Typically, digital image sensor ISOs range from 100 (fairly slow) to 3200 or higher (very fast).

Some digital cameras have more than one ISO rating. In low-light situations, the sensor’s ISO can be increased by amplifying the image sensor’s signal (increasing its gain). Some cameras even increase the gain automatically. This not only increases the sensor’s sensitivity, but, unfortunately, also increases the noise or “grain,” thus, generating images that are contaminated with random noise effects.

2 Sensor Noise Types

Noise can be summarized as the visible effects of an electronic error (or interference) in the final image from a digital camera. Noise is a function of how well the image sensor and digital signal processing systems inside the digital camera are prone to and can cope with or remove these errors (or interference). Noise significantly degrades the image quality and increases the difficulty in discriminating fine details in the image. It also complicates further image processing, such as image segmentation and edge detection. The type of high-ISO sensor noise produced by a typical digital camera CCD imaging sensor can be modeled as an additive white Gaussian distribution with zero mean and a variance (noise power) proportionate to the amount of amplification applied to the image sensor’s signal to boost its gain.^{1–3}

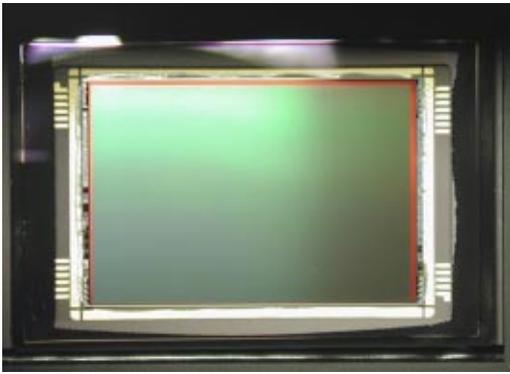


Fig. 1 The CCD imaging sensor chip used in the “Kodak Professional DCS720x” digital camera. (courtesy <http://www.dpreview.com/>).

Visible noise in a digital image is often affected by temperature (high worse, low better) and ISO sensitivity (high worse, low better). Some cameras exhibit almost no noise and some a lot and all the time. It has certainly been the challenge of digital camera developers to reduce noise and produce a “cleaner” image, and indeed some recent digital cameras are improving this situation greatly, allowing for higher and higher ISOs to be used without too much noise. In general, image artifacts produced by digital cameras can be divided into three types.

- *Stuck-pixel noise*: also known as impulse-type noise is created by many digital cameras and is caused by long exposure time of the CCD elements in dim-lighting conditions when a bright image is required without the use of the flash light. During the exposure, some CCD cells become saturated and are stuck at a bright color, which show up in the acquired image as bright impulse noise pixels [see Fig. 2(a)]. Removing this type of noise is usually not a difficult task but comes at the cost of blurring the other uncorrupted pixels.
- *Blue-channel noise*: is a common problem in digital photographs, especially those created by high-ISO professional news and sports digital cameras. A typical digital camera sensor [CCD or complementary metal–oxide–semiconductor (CMOS)] is more sensitive to certain primary colors than others (often sensors are

less sensitive to blue light) and so to compensate, these channels are amplified more than the others [see Fig. 2(d)]. This noise has hampered the acceptance of digital cameras for quality reasons, and it has limited their use for some techniques such as automasking based on chrominance (e.g., “blue screen” backgrounds).

- *JPEG artifacts*: also known as JPEG blocking artifacts, are due to the nature of the 8-by-8 block-size used by the JPEG compression standard. High compression ratios result in images with blockiness in the blue and red channels. These blocks are especially obvious in the flat areas of an image. In high detail areas, artifacts called “mosquito noise” become noticeable. This term comes from the ripple effect that mosquitoes make when their legs touch water.

The sensor noise filtering techniques that will be described shortly present a solution to these types of artifacts that are common in most color images acquired by digital cameras.

3 Background of Noise Reduction Techniques

The primary concern in digital photography is the visual fidelity of the acquired images. What professional photographers demand from a digital camera is fast and precise image acquisition (high-sensitivity in low-light conditions and exact-moment capture) coupled with the best visual results. Previous noise reduction techniques available in the literature do not take into account the physics of the CCD photo-capture element used in a majority of video and digital still cameras today. The CCD image sensor tends to be highly sensitive to the green light frequencies and less sensitive to the blue and red light waves. The CCD hardware controller is, thus, tuned to increase the signal gain of the blue CCD elements more than the green elements. In normal and good lighting conditions there are no visible effects due to this difference in signal gain, but in low light conditions, where the CCD signal gain is increased more for the less sensitive color channels, this produces high-frequency noise that contaminates the blue channel, and, to a lesser extent, the red channel, more severely than the green channel. In general the chrominance channels of the acquired images will be more severely affected by this

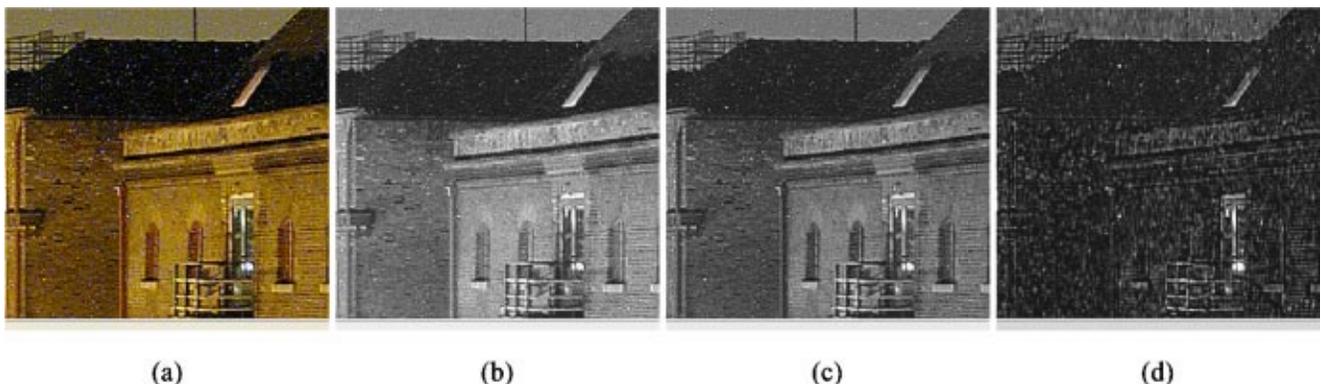


Fig. 2 (a) Long exposure stuck pixel noise (courtesy <http://www.dpreview.com/>), (b) red channel, (c) green channel, and (d) blue channel showing excessive noise.

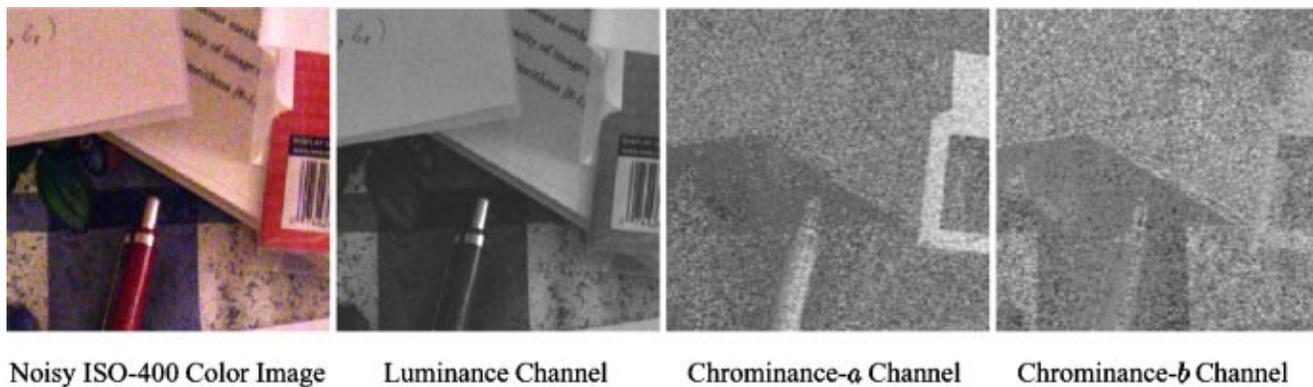


Fig. 3 An image acquired in low light at an ISO-400 setting with a *Minolta Dimage* digital camera, then separated into its component channels using the $L^*a^*b^*$ color space to show the severity of noise in the chrominance a and b channels as compared to the luminance channel due to the limitations of the CCD image sensor.

noise than the luminance channel of the image as shown in Fig. 3. The resulting effect is the visibility of random noise artifacts in the acquired image that differs in severity from acceptable (at low-ISO settings $< \text{ISO } 400$) to completely contaminating the picture (at very high ISO settings $> \text{ISO } 2000$) such that it becomes visually unacceptable. More details on separating a color image into luminance (brightness) and chrominance (color) channels will be presented in Sec. 4.

Statistical characteristics of images are of fundamental importance in many areas of image processing. Incorporation of *a priori* statistical knowledge of spatial correlation in an image, in essence, can lead to considerable improvement in many image processing algorithms. For noise filtering, the well-known Wiener filter for minimum mean-squared error (MMSE) estimation is derived from a measure or an estimate of the power spectrum of the image, as well as the transfer function of the spatial degradation phenomenon and the noise power spectrum.⁴ Unfortunately, the Wiener filter is designed under the assumption of wide-sense stationary signal and noise. Although the stationarity assumption for additive, zero-mean, white Gaussian noise is valid for most cases, it is not reasonable for most realistic images, apart from the uninteresting case of uniformly gray image fields. What this means in the case of the Wiener filter is that we will experience uniform filtering throughout the image, with no allowance for changes between edges and flat regions, resulting in unacceptable blurring of high-frequency detail across edges and inadequate filtering of noise in relatively flat areas.

Noise reduction filters have been designed in the past with this stationarity assumption. These have the effect of removing noise at the expense of signal structure. Examples such as the fixed-window Wiener filter and the fixed-window mean and median filters have been the standard in noise smoothing for the past 2 decades.⁴⁻⁶ These filters typically smooth out the noise, but destroy the high-frequency structure of the image in the process. This is mainly due to the fact that these filters deal with the fixed-window region as having sample points that are stationary (belonging to the same statistical ensemble). For natural scenes, any given part of the image generally differs sufficiently from the other parts so that the stationarity assumption

over the entire image or even inside a fixed-window region is not generally valid. Newer adaptive Wiener filtering techniques that take into account the nonstationarity nature of most realistic images have been used as an alternative to preserve signal structure as much as possible.⁷ Many, however, do this at the expense of proper noise reduction, where the high-frequency areas will be insufficiently filtered, which will result in a large amount of high-amplitude noise remaining around edges in the image. Another shortcoming is the failure of these filters to remove stuck-pixel (impulse-type) noise that appears in the acquired images due to long CCD exposure times in dim light. A fixed-window median filter will remove this type of impulse noise but will also alter important signal structure due to the same assumption that the image samples in the fixed-window can be modeled by a stationary random field, which is not valid for a fixed-window that cannot inherently differentiate between edge and flat image regions.⁸

Another shortcoming with many noise filtering techniques that deal with color digital images is the application of the same filter evenly to the three color channels (R,G,B)^{9,10} with the assumption that the sensor noise is equally distributed among the three color channels, which is an erroneous assumption as explained at the start of this section. To the author's knowledge, the issue of high-ISO noise reduction for digital cameras has not received much attention in the literature. Although current work in the literature on adaptive noise reduction filters (such as Smolka *et al.*,¹¹ Eng and Ma¹²) may be used to reduce high-ISO digital camera noise, these filters have been developed without the specific needs of professional digital cameras and as such do not take into consideration the different types of noise degradations which are generated by these digital cameras as mentioned previously. The result is either insufficient noise reduction in the chrominance channels, or too much smoothing in the luminance channel of the filtered image. Also, many filters only deal with gray-level images^{6,8,13-17} and, as such, are not very effective for the type of images acquired by the digital cameras discussed in this work. We will compare one of the commonly used adaptive spatial noise reduction filters, namely the adaptive local statistic MMSE filter, with our work to show the effectiveness of our technique in producing visually superior

filtered images that can directly be used for further analysis, such as edge detection and image understanding.

Digital camera manufacturers are only now beginning to realize the importance of incorporating noise reduction filters in the hardware image acquisition pipeline for their digital cameras. To the author's knowledge, the only known digital camera to actually attempt to incorporate a noise reduction algorithm is the Kodak Professional DCS720x digital camera.¹⁸ Kodak rates the camera as "calibrated" up to ISO 4000 and capable of ISO 6400. This means that shooting in extremely low light/high shutter speed conditions is possible with this camera, but at the expense of increased noise in the acquired images. With noise reduction activated the images are acquired with much less noise and are more visually pleasing. Images from this camera are used as a comparison with the techniques described in this paper.

In the next sections, a new adaptive technique is described that is highly tuned to produce visually pleasing filtered color digital images that have been acquired using digital sensor-based cameras. This new technique differs from previous filtering methods in that it is geared towards the type of color images obtained from digital cameras, and thus takes into account the physical limitations of the CCD and the specific types of CCD sensor noise produced.

4 Color Spaces

Before going into the details of the new color filtering methods, it is important to give a brief background of the most popular color spaces used in separating color images into their component color channels before filtering each channel.

A color space is a model for representing color in terms of intensity values. It defines a one-, two-, three-, or four-dimensional space whose dimensions, or components, represent intensity values. A color component is also referred to as a color channel. For example, RGB space is a three-dimensional color space whose components are the red, green, and blue intensities that make up a given color. Color spaces can be divided into two general categories; device dependent and device independent color spaces.

- **Device dependent color spaces:** These include the family of RGB Spaces. The RGB space is a three-dimensional color space whose components are the red, green, and blue intensities that make up a given color. Most CCD and CMOS based digital camera imaging sensors use the RGB color space by reading the amounts of red, green, and blue light reflected from the scene that fall on the CCD elements and then convert those amounts into digital values. These values are device dependent and one CCD may produce a different RGB value from another depending on how it is manufactured. HSV space and HLS space are transformations of RGB space that can describe colors in terms more natural to an artist. The name HSV stands for hue, saturation, and value, and HLS stands for hue, lightness, and saturation. The CMY color space is sometimes used in CCD and CMOS image sensors. The name CMY refers to cyan, magenta, and

yellow, which are the three primary colors in this color space, and red, green, and blue are the three secondaries.

- **Device independent color spaces:** Some color spaces can express color in a device-independent way. Whereas RGB colors vary with CCD and CMOS sensor hardware characteristics, device-independent colors are meant to be true representations of colors as perceived by the human eye. These color representations, called device-independent color spaces, result from work carried out in 1931 by the Commission Internationale d'Eclairage (CIE) and for that reason are also called CIE-based color spaces. The CIE created a set of color spaces that specify color in terms of human perception. It then developed algorithms to derive three imaginary primary constituents of color, namely X , Y , and Z , that can be combined at different levels to produce all the color the human eye can perceive. The resulting color model, and other CIE color models, form the basis for all color management systems. Although the RGB and CMY values differ from device to device, human perception of color remains consistent across devices. Colors can be specified in the CIE-based color spaces in a way that is independent of the characteristics of any particular imaging device. The goal of this standard is for a given CIE-based color specification to produce consistent results on different devices, up to the limitations of each device.¹⁹

One problem with representing colors using the XYZ color space is that it is perceptually nonlinear: it is not possible to accurately evaluate the perceptual closeness of colors based on their relative positions in XYZ space. Colors that are close together in XYZ space may seem very different to observers, and colors that seem very similar to observers may be widely separated in XYZ space. $L^*a^*b^*$ space is a nonlinear transformation of XYZ space to create a perceptually linear color space designed to match perceived color difference with quantitative distance in color space.^{20,21}

As stated earlier, the CCD sensor of a typical digital camera is less sensitive to the blue and red channels and this causes amplified noise artifacts in the chromatic channels in low light or at high-ISO settings. Moreover, there seems to be general agreement that spatial resolution is markedly lower in chromatic channels than in the achromatic one (see Fig. 3), hence, high-frequency information, i.e., edges, come mainly from this achromatic channel.²² Another important consideration is that, in order to avoid chromatic artifacts in the filtered image, a nonlinear operator cannot be applied to each RGB component separately.²² These two considerations and experimental results suggest that a color model which should separate luminance from chrominance is suitable. We, thus, choose to separate our acquired images using the $L^*a^*b^*$ color space because of its merits stated earlier. The $L^*a^*b^*$ color space separates the RGB image into a luminance channel L , and two chrominance channels (a,b) . This allows us to use different filter parameters specifically tuned for each channel. In general the luminance channels suffer less noise artifacts

than the (a,b) chrominance channels. We, therefore, take this into consideration when filtering each channel, by allowing more smoothing in the (a,b) channels to correct for color artifacts, while passing more high frequency in the filtered luminance channel. This will be further emphasized when presenting the experimental results in a later section.

5 Adaptive-Window Signal Equalization Hybrid Filter

The quantity of light falling on an image sensor array (e.g., CCD array), is a real valued function $q(x,y)$ of two real variables x and y . An image is typically a degraded measurement of this function, where degradations may be divided into two categories, those that act on the domain (x,y) and those that act on the range q . Sampling, aliasing, and blurring act on the domain, while noise (including quantization noise) and the nonlinear response function of the camera act on the range.²³ We are concerned with the latter type of camera sensor degradations.

Digital camera sensor noise reduction is the process of removing unwanted noise from a digital image. It falls into two main categories, reduction or removal of noise from high-ISO images including JPEG compression artifacts, and reduction or removal of noise from long-exposure images (with “stuck pixels”). In this section a detailed description of the adaptive hybrid filter for sensor noise removal is presented with experimental results showing its performance in comparison with other standard noise reduction filters.

5.1 Hybrid Mean and Adaptive Center Weighted Median Filter

The median filter is a class of *order-statistic filters* where filter statistics are derived from ordering (ranking) the elements of a set rather than computing means, etc. The median filter is a nonlinear neighborhood operation, similar to convolution, except that the calculation is not a weighted sum. Instead, the pixels in the neighborhood are ranked in the order of their gray levels, and the midvalue of the group is stored in the output pixel. In probability theory, the median M , of a random variable x , is the value for which the probability of the outcome $x < M$ is 0.5.⁶ Median filtering is normally a slower process than convolution, due to the requirement for sorting all the pixels in each neighborhood by value. There are, however, algorithms that speed up the process.^{24,25} The median filter is popular because of its demonstrated ability to reduce random impulsive noise without blurring edges as much as a comparable linear low-pass filter. However, it often fails to perform as well as linear filters in providing sufficient smoothing of nonimpulsive noise components such as additive Gaussian noise. In order to achieve various distributed noise removal, as well as detail preservation, it is often necessary to combine linear and non-linear operations.^{26–29} In this section we introduce a hybrid filter combining the best of both worlds; proper smoothing in flat regions and detail preservation in busy regions of the image.

One of the main disadvantages with the basic median filter is that it is location-invariant in nature, and thus also tends to alter the pixels not disturbed by noise. The center weighted median filter (CWMF) was developed to address

this limitation in the basic median filter.^{30,31} This filter will give the pixel at the center of the window more weight (>1) than the other pixels in the window before determining the median. This has the effect of preferentially preserving that pixel’s value so both fine detail and noise are more preserved. In the extreme, one could make it so that the center pixel has a weight equal to the entire weight of the rest of the window, in which case the value of the center pixel is assured of being the output of the median operation. This is the identity filter, where the output is equal to the input. In general, a CWMF can be varied over the range from the median filter to the identity filter by varying the central weight. This corresponds to the range from strong noise and detail removal (basic median filtering) to none (identity filtering). In the original CWMF implementation, the central weight is constant over the entire image.

In this paper, we make use of the CWMF concept and implement it as an adaptive CWMF (ACWMF) by varying the central weight based on signal and noise estimates inside an adaptive window framework, which will be described in detail in the next section.

In formulating our center weighted median-based filter we use an image model with additive noise as follows:

$$y(k,l) = x(k,l) + n(k,l), \quad (1)$$

where $k \in [0, M-1]$, and $l \in [0, N-1]$ for an $M \times N$ sized image. $n(k,l)$ is a zero-mean additive white Gaussian noise random variable, of variance σ_n^2 , and uncorrelated to the ideal image $x(k,l)$, which is assumed to be of zero mean and variance σ_x^2 , and $y(k,l)$ is the noise-corrupted input image.

For the purpose of the following analysis, we assume that both $x(k,l)$ and $n(k,l)$ are ergodic random variables. The implication of this assumption is that although we do not have *a priori* knowledge of the signal and noise statistical variance and mean, we can still capture samples of $x(k,l)$ and $n(k,l)$ and determine their variance and mean, which are, in turn, representative of their respective ensembles. It should also be noted that although the noise variance, σ_n^2 , is not known *a priori*, it is easily estimated from a window in a flat area of the degraded image $y(k,l)$.³²

We begin by setting an objective criterion of optimality for deriving the central weight at each pixel location. We use a similar criterion to that used in deriving the power spectrum equalization filter,³³ by seeking a linear estimate, $\hat{x}(k,l)$, such that the signal variance of the estimate is equal to the variance of the ideal image $x(k,l)$. Assuming this estimate is of the form

$$\hat{x}(k,l) = \alpha \cdot y(k,l), \quad (2)$$

we can express our criterion as

$$\sigma_x^2 = E\{\hat{x}^2\} = E\{(\alpha \cdot y)^2\}, \quad (3)$$

where $E\{\cdot\}$ is the expectation operator. In general, the acquired images have a nonzero mean, and we can account for this by subtracting the mean of each image from both random variables of Eq. (2). For zero mean noise, the a

posteriori sample mean (local mean inside the adaptive window) of the degraded pixel $y(k,l)$, denoted by m_y , is equal to the *a priori* sample mean of the ideal pixel $x(k,l)$. After dropping the (k,l) notation for readability, we have

$$\hat{x} - m_y = \alpha \cdot (y - m_y), \quad (4)$$

and we can write our criterion after accounting for the mean as follows:

$$\begin{aligned} \sigma_x^2 &= E\{(\hat{x} - m_y)^2\} = E\{\alpha \cdot (y - m_y)\}^2 \\ &= \alpha^2 \cdot \sigma_y^2 = \alpha^2 \cdot (\sigma_x^2 + \sigma_n^2). \end{aligned} \quad (5)$$

Therefore, the signal equalization estimator α becomes

$$\alpha = \sqrt{\frac{\sigma_x^2}{\sigma_x^2 + \sigma_n^2}}. \quad (6)$$

If the number of pixels in the adaptive window, of size $L_x \times L_y$, is $(L_x \cdot L_y)$, then the central weight for the pixel under analysis at pixel position (k,l) is given as

$$C_w = \alpha \cdot (L_x \cdot L_y - 1) + 1. \quad (7)$$

This central weight can be used to give the value of the pixel at (k,l) more weight than the other pixels in the adaptive window before determining the median; i.e., we count it as if it were C_w pixels rather than just one pixel. Thus, $\alpha=0$ is the basic median filter, while $\alpha=1$ is the identity filter, and $0 \leq \alpha \leq 1$.

Substituting α from Eq. (6) in Eq. (4), we obtain an equation for the signal equalization mean filter which can estimate the ideal image as

$$\hat{x} = m_y + \sqrt{\frac{\sigma_x^2}{\sigma_x^2 + \sigma_n^2}} \cdot (y - m_y). \quad (8)$$

We, thus, have two ways (filters) for estimating the ideal image. The nonlinear, order-statistic-type ACWMF with a central weight given by Eq. (7), and the linear signal equalization mean filter of Eq. (8). We propose to use a hybrid combination of linear and nonlinear operations within the adaptive window framework described in the next section. The new hybrid filter is given by

$$H(k,l) = \begin{cases} \text{ACWMF}(k,l) & \text{if } \alpha > \eta \\ m_y(k,l) & \text{otherwise} \end{cases}. \quad (9)$$

Here, η is an empirically determined threshold for choosing between the ACWMF and the signal mean m_y . We use the value $\eta=0.8$ for best results. Thus, when $\alpha > \eta$ indicating high signal activity in edge regions, the ACWMF is preferred over the mean for proper noise removal with minimal blurring. Otherwise, in flat regions of the image, α is small, and the local mean m_y inside the adaptive window is selected for smoothing.

Replacing the sample mean m_y of Eq. (8) with the hybrid mean-ACWMF filter (HM-ACWMF), H , we obtain the final equation for the ideal signal equalized estimate as

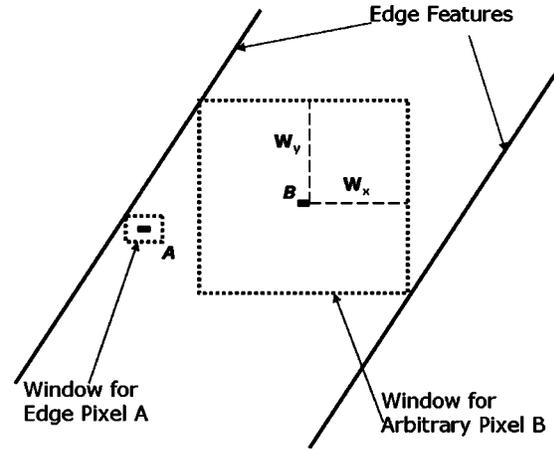


Fig. 4 An adaptive rectangular window shown for an edge pixel A and an arbitrary pixel B situated away from the edge.

$$\hat{x}(k,l) = \alpha \cdot y(k,l) + (1 - \alpha) \cdot H(k,l). \quad (10)$$

Since H depends on local (sample) statistics such as $\sigma_x^2 = \sigma_y^2 - \sigma_n^2$ and m_y , proper values of H will depend on appropriate adaptive window dimensions that do not cross-over object boundaries in order for the stationarity assumption to hold inside the adaptive window region. The following section details a new optimal-size adaptive window framework that is utilized for the HM-ACWMF described in this section.

5.2 Optimal-Size Windowing

In developing an effective adaptive window to account for the nonstationarity of images, it is important for this analysis window to have the maximum size possible at each pixel position without crossing over image structure and edges. The reason is that the more the number of spatially correlated signal samples are available in the analysis window, the more accurate their statistical characteristics can be estimated.³⁴ Previous adaptive windowing techniques are not optimal in this respect as they usually vary the window size in the right-most upper quadrant only; namely the positive x axis and the positive y axis, and simply duplicate these values in the negative axis. Thus, the window grows in size symmetrically, increasing to the maximum (user-preset) allowable size in the middle of a flat region of the image and decreasing to the minimum size near the edges.^{7,17} Figure 4 shows two different situations for this type of an adaptive window where the window is symmetric and decreases to a minimum size for the edge pixel A, thus causing inaccurate estimates of signal statistics. The window will increase to its maximum size for pixels (e.g., pixel B) further away from the edge region where more accurate statistical estimates can be obtained.

In designing the adaptive window for the signal equalization filter, the optimality of window size at every pixel position has been taken into consideration. Instead of maintaining window symmetry, we are more concerned with the window size near the edges of image structure (object boundaries). The new adaptive window is structured such

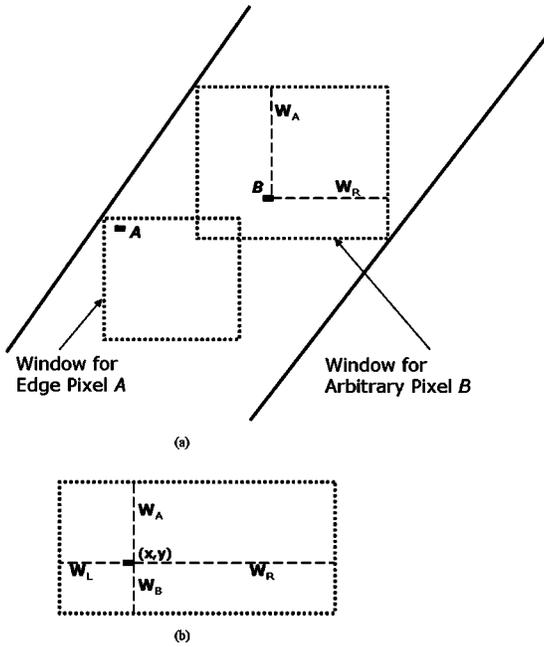


Fig. 5 (a) The optimal-size adaptive window shown for an edge pixel *A* and an inner pixel *B* situated away from the edge. (b) The structure of the window is shown, where each of the four quadrants W_R , W_L , W_A , W_B is separately adaptive based on the signal strength inside the window at maximum size.

that all four quadrants of the window are adaptive, and vary in size based on an estimate of the signal activity in the respective window quadrant.

Figure 5(a) shows the same example pixels (*A* and *B*) as those in Fig. 4. The figure illustrates the big difference between the two adaptive window examples. In our case, the adaptive window grows to its maximum allowable size even for the edge pixel *A*. The two methods are comparable, though, for the inner pixel *B* as shown in Fig. 5(a). We, thus, have an adaptive-size window that is optimal for both edge pixels as well as pixels that belong to flat regions of the image.

Referring to Fig. 5(b), the window is started with its four quadrants set to the maximum (user-defined) size at every pixel position (k,l) . A measure of the uncorrupted signal activity in this window at pixel (k,l) is given by the local signal variance $\sigma_x^2 = \sigma_y^2 - \sigma_n^2$. This signal variance in the current window is then compared with the noise variance σ_n^2 , and if $\sigma_x^2 > \zeta \cdot \sigma_n^2$, it is assumed that the window has identified a region in the ideal image with significant structural characteristics such as an edge. This is contrary to the desired characteristics of the adaptive window. The adaptive window is to be formed such that it includes relatively uniform structures in the ideal image, so that the primary source of variance in the window is the additive noise. Therefore, the window is reset to its minimum size, a threshold T is set to the value of the estimated signal activity, $T = \sigma_x^2$, and the window size is increased iteratively one quadrant at a time. A new signal activity is estimated in the first quadrant (W_R) and as long as it is less than the threshold T the axis W_R is incremented to $W_R + 1$. The signal activity is then re-estimated in the new window

($W_R + 1, W_L, W_A, W_B$) and compared to the threshold T to either increment W_R again or stop at the current length. This continues iteratively until either the signal activity in the new window rises above the threshold T , indicating closeness to an edge, or the current window axis W_R reaches the maximum (user preset) value. This is repeated for the other three axes W_L, W_A, W_B in a similar manner. On the other hand, if $\sigma_x^2 < \zeta \cdot \sigma_n^2$, then it is assumed that the window is in an edge-free (flat) region of the image and the window will stay at the maximum dimension for the current pixel. This will result in maximum smoothing of noise and JPEG artifacts. Here ζ is a (user defined) weighting constant that affects the amount of smoothing. Large values of ζ will cause more noise to be filtered, but may also result in smaller signal structure being blurred. This value is empirically selected by the user to give the best results, and is image dependent.

The described optimal-size adaptive window framework allows for bigger window sizes to be used safely without the danger of excessive blurring of edges as the size grows optimally large even near edges but does not cross over the edge, as Fig. 5(a) clearly illustrates. It should also be noted that the HM-ACWMF formulation of Eq. (10) is applied to each channel of the $L \times a \times b$ space separately while varying the filter parameters (upper-limit for the adaptive window size, ζ , and estimated σ_n^2) depending on the type of channel being filtered. For example the estimated noise variance of the L channel is usually less than that of the a channel or the b channel. Therefore, a smaller adaptive window is sufficient and a smaller value for ζ will produce the best results. These values are usually increased for the chrominance (a, b) channels due to the increase in sensor noise in these channels.

In the remainder of this section, experimental results are presented with various types of digital images corrupted by high-ISO sensor noise, including the unavoidable JPEG compression artifacts, at varying degrees of severity.

5.3 Experimental Results for the HM-ACWMF

In this section, experimental results are presented to show the performance of our HM-ACWMF noise filter when applied to ISO-noise corrupted digital images, and to compare its performance with two other commonly used filters; the fixed window median filter and the adaptive local statistic MMSE filter given by¹⁶

$$\text{MMSE} = y(k,l) + \frac{\sigma_n^2}{\sigma_y^2} \cdot [y(k,l) - m_y(k,l)]. \quad (11)$$

In evaluating the performance of individual filters it is important to take into consideration both the analytical performance of the filter as well as the visual quality of the estimated images generated by the filter. The well known MSE metric calculates the amount of difference between an ideal image and its estimate, and has been widely used for measuring the performance of various filters.⁴ The use of the MSE metric for measuring filter performance can be justified for single-channel filters that process gray-scale images, but for multichannel (color) image processing, a compound MSE metric would be more appropriate to mea-

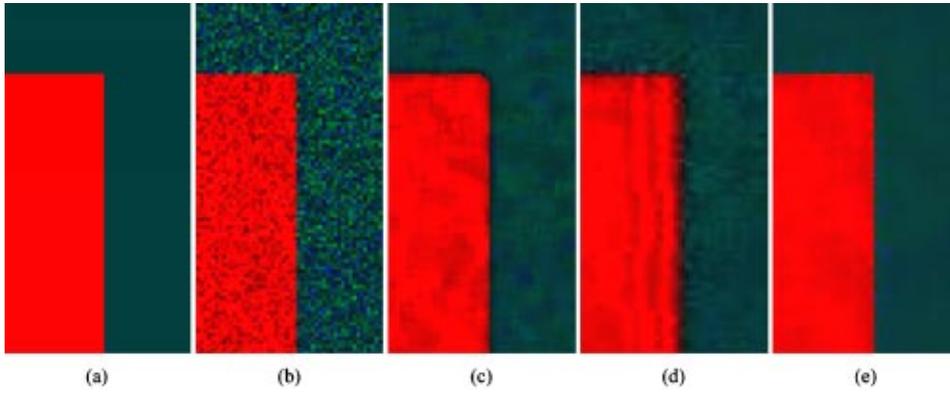


Fig. 6 (a) A simulated noise-free edge color image. (b) The same image corrupted by white Gaussian noise with (L, a, b) -channel variances of $\sigma_n^2 = (20, 70, 100)$ in the $L^*a^*b^*$ color space, respectively. (c) Median filtered image with 5×5 mask size and $MSE_{Lab} = (1.8, 8.3, 12.9)$. (d) MMSE filtered image with 7×7 kernel mask and $MSE_{Lab} = (2.5, 14.4, 16.5)$. (e) HM-ACWMF filtered image with 11×11 maximum adaptive window size, $\zeta = 0.1$, and $MSE_{Lab} = (0.4, 7.9, 11.3)$.

sure the difference between individual channels of a multi-channel image, where the MSE for an arbitrary x channel can be given by

$$MSE_x = \frac{1}{M \cdot N} \sum_{i=0}^{M \cdot N - 1} (x_i - \hat{x}_i)^2, \quad (12)$$

for an $M \times N$ size image, with x representing the ideal image channel, and \hat{x} the estimated image channel. For RGB images, we can compute a separate MSE for the R, G, and B channels. The only shortcoming in a RGB MSE metric is that it is not ideal for tracking visual quality in an estimated image, because, as explained in Sec. 4, the RGB color space is device dependent and does not represent true colors perceived by the human eye. In comparison, the $L^*a^*b^*$ color space is a device-independent color space, that is a true representation of colors as perceived by the human eye. Using a $MSE_{Lab} = (MSE_L, MSE_a, MSE_b)$ metric based on the $L^*a^*b^*$ color space makes more sense and should be capable of emphasizing the strength of the HM-ACWMF as compared to the other filters used in the comparison. It is also important to note that the filters used for comparison are standard filters reported in the literature and they simply apply the same filter parameters evenly to the individual R,G,B channels in the RGB color space. They thus assume that the amount of noise is evenly distributed among the three RGB color channels which is an invalid assumption. As explained earlier, the chrominance channels are more severely affected by sensor noise than the luminance channel. A filter that works in the $L^*a^*b^*$ color space while properly tuning the filter parameters to deal with noise on a per-channel basis (HM-ACWMF) is expected to generate image estimates of higher visual quality, which can be evaluated by comparing the chrominance MSE_a and MSE_b metrics of each filter.

We start by showing results of the performance of the adaptive window with the hybrid filter near edges using a synthetic image. Figure 6(a) shows an enlarged simulated edge with two different colors. This image was corrupted by a Gaussian random variable of zero mean and variance

$\sigma_n^2 = (20, 70, 100)$ in the (L, a, b) channels of the $L^*a^*b^*$ color space respectively, as shown in Fig. 6(b). The next two Figs. 6(c) and 6(d) show the median and MMSE filtered images, respectively. It is clear that noise artifacts remain in the filtered images indicating unsatisfactory results. The HM-ACWMF filtered image is shown in Fig. 6(e) with a $L^*a^*b^*$ mean square error $MSE_{Lab} = (0.4, 7.9, 11.3)$ indicating a reduction in the noise levels from $(20, 70, 100)$ as is apparent from the image. The important aspect in this image is the sharpness of the edges, which proves that the optimal-size adaptive windowing framework developed in this paper performs as expected.

We now present some experimental results with real images acquired by CCD-based digital cameras at high-ISO settings. Figure 7(a) shows a portion of a larger image used as a reference for measuring the performance of the noise filters and acquired at an ISO setting equivalent to ISO-50. Figure 7(b) shows the same portion of the image acquired at an ISO-400 setting and stored in the JPEG compressed format. It is clear that the image suffers from high-ISO sensor noise. The MSE_{Lab} noise variance was estimated at $\sigma_n^2 = (3.5, 3.5, 3.3)$. Figure 7(g) shows the gradient image of (b) to emphasize the JPEG blocking artifacts. Figure 7(c) shows the image filtered using the 3×3 median filter. It is clear both visually and in terms of the $L^*a^*b^*$ mean squared error [MSE_{Lab} has increased to $(33.8, 4.2, 8.7)$ from $(3.5, 3.5, 3.3)$] that the filter has completely failed to adequately remove the sensor noise as is apparent in the flat regions. Figure 7(d) shows the MMSE filtered image with a kernel mask size of 9×9 . The luminance MSE_L has been reduced to 2.9 from 3.5, but the chrominance MSE_{ab} , which measures the visual quality, has increased to $(3.9, 3.6)$ indicating that the filter has failed to completely remove the ISO noise and JPEG artifacts from flat regions, as is apparent from the many artifacts left behind [apparent in the gradient image of Fig. 7(i)], which greatly degrade its visual appearance.

Figure 7(e) is the HM-ACWMF filtered image. Filter parameters used were $\zeta = 0.1, \eta = 0.8$. The noise variance, MSE_{Lab} of $(1.6, 3.5, 3.0)$, has been significantly reduced in

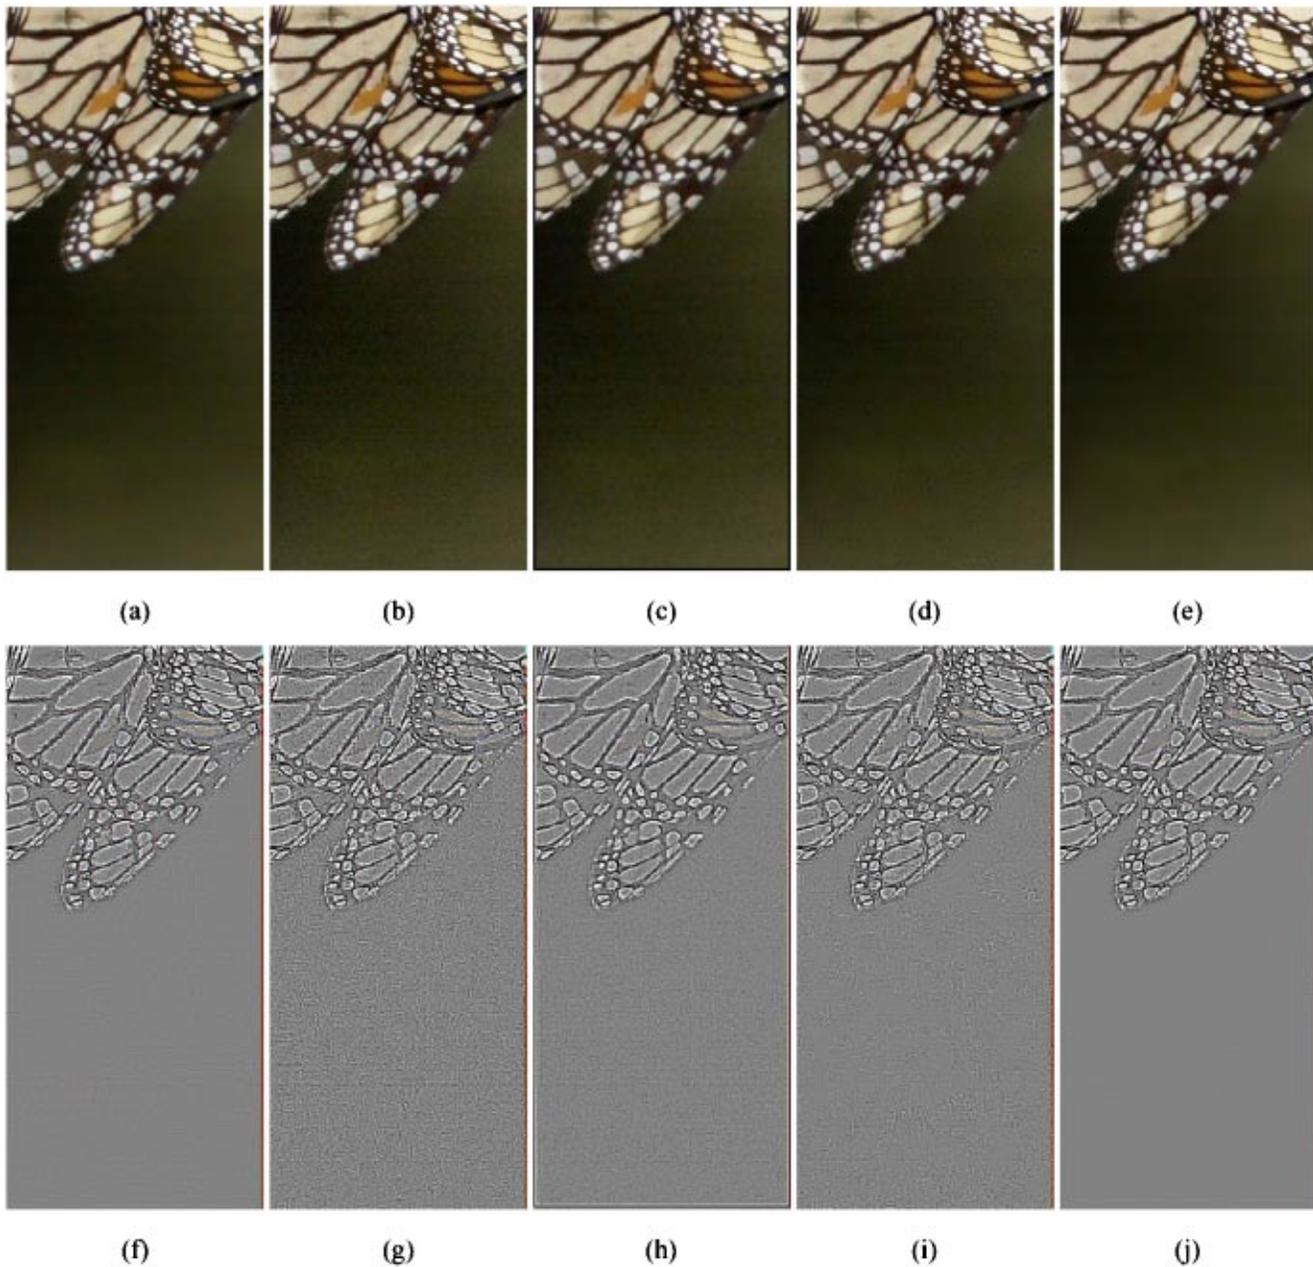


Fig. 7 (a) An image showing part of a butterfly's wing taken at ISO-50 as a reference (noise-free) image. (b) The same image taken at ISO-400 showing two types of noise; ISO-noise and JPEG blocking [$\sigma_n^2 = (3.5, 3.5, 3.3)$]. (c) Median filtered image with 3×3 mask [$MSE_{Lab} = (33.8, 4.2, 8.7)$]. (d) MMSE filtered image with 9×9 kernel mask [$MSE_{Lab} = (2.9, 3.9, 3.6)$]. (e) HM-ACWMF filtered image (maximum adaptive window size of 11×11 , $\zeta = 0.1$), showing sharp edges and uniform smoothing in flat regions [$MSE_{Lab} = (1.6, 3.5, 3.0)$]. (f)–(j) High-pass filtered gradient equivalent of images in (a)–(e). (g) shows the JPEG blocking artifacts clearly in the flat region. (j) shows complete removal of JPEG artifact noise from the flat region, as well as ISO-noise from busy regions, which accounts for the superior quality compared to the other filters.

the luminance and chrominance- b channels, and remained the same for the chrominance- a channel indicating an improved overall visual quality. The JPEG artifacts that severely affected the degraded image have also been completely removed from the flat regions. The optimally adaptive window was allowed to increase to a maximum size of 11×11 , which accounts for the clean look in the flat region. It is clear that this image was acquired in daylight

and, as such, there was no need for the imaging sensor to overamplify the signal gain for the chromatic channels at this ISO setting (ISO-400). All three channels [luminance, and chrominance (a, b) of the $L^*a^*b^*$ space] were, thus, filtered using the same filter settings due to the relatively evenly distributed small amount of ISO noise in the three channels of the degraded image. Figure 7(j) shows the gradient image of Fig. 7(e) with a very smooth flat region and

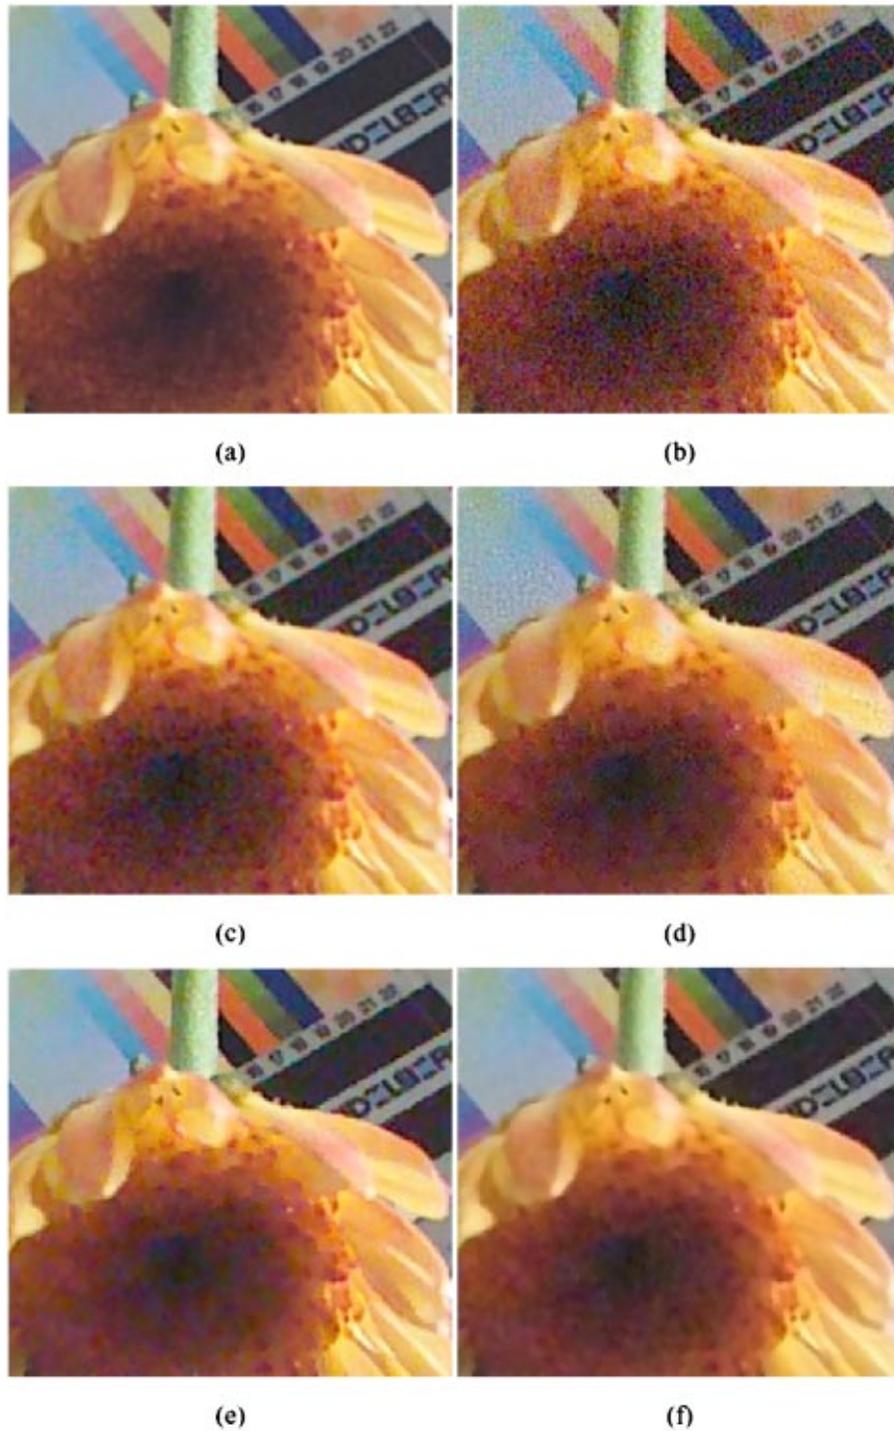


Fig. 8 (a) Flower reference image at ISO-400 (courtesy <http://www.dpreview.com/>). (b) The severely corrupted ISO-6400 image, with a measured MSE_{Lab} noise variance of (21.4, 30.4, 71.3). (c) Median filtered image with 3×3 mask [$MSE_{Lab} = (45.9, 28.5, 69.9)$]. (d) MMSE filtered image with 7×7 kernel mask [$MSE_{Lab} = (14.6, 19.7, 48.6)$]. (e) Kodak's proprietary ISO-noise filter built into the camera firmware [$MSE_{Lab} = (9.8, 20.0, 57.3)$]. (f) HM-ACWMF nonlinear filtered flower image using 5×5 maximum adaptive window size and $\zeta = 1$ for the luminance channel [MSE_{Lab} has dropped to (8.7, 14.6, 37.3)].

completely cleared JPEG and ISO noise artifacts.

The butterfly image was a relatively difficult test for the adaptive window HM-ACWMF because of the highly busy regions in the butterfly wings, yet, the ISO-400 noise

level was on the low side. Next, we show a more complicated image acquired by the Kodak Professional DCS720x digital camera at its highest ISO setting of 6400. We compare our technique with Kodak's own built-in noise filter.

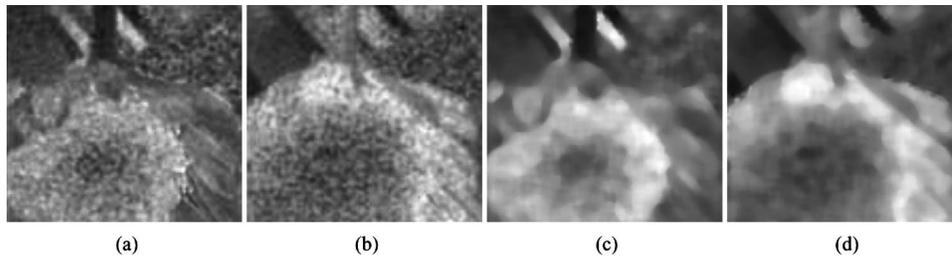


Fig. 9 (a) The noise corrupted a channel, showing severe noise that causes color artifacts as in Figs. 8(b), 8(c), 8(d), 8(e). (b) The noise corrupted b channel. (c) HM-ACWMF filtered a channel using 13×13 maximum adaptive window size and $\zeta=4$ for proper flattening of the color artifacts. (d) HM-ACWMF filtered b channel using the same filter parameters as for filtering the a channel.

We take an ISO-400 copy of the image (also taken by the DCS720x) as an ideal noise-free reference for performance comparison. Figure 8(a) shows the ISO-400 reference image of a flower. Figure 8(b) shows the same image acquired in a well-lit environment at ISO-6400 showing severe ISO-noise that affects all color channels as apparent from the nonuniform color artifacts that appear in the noisy image, and absent from the reference image in (a). The estimated σ_n^2 noise variance was (21.4,30.4,71.3).

Figure 8(c) is the median filtered image with a mask of size 3×3 . The MSE_{Lab} between the median filtered image and the reference image was measured at (45.9,28.5,69.9). It is clear that the basic median filter is inadequate, both visually as apparent from the severe blurring and the high chrominance MSE_{ab} values, and analytically from the large increase in the MSE_L value. The next image of Fig. 8(d) is the MMSE filtered image using a kernel size of 7×7 . The MSE_{Lab} measure was (14.6,19.7,48.6), a slight improve-

ment over the ISO-6400 noisy image, but still lacking the visual fidelity that professional digital photography demands. Kodak's own noise filtered image is shown in Fig. 8(e). The measured MSE_{Lab} measure was (9.8,20.0,57.3). One observation here is that the luminance MSE_L gives a measure of the amount of noise reduction in the filtered image as apparent from the lower MSE_L value for the Kodak image as compared to both the MSE_L values for the median and the MMSE filtered images (9.8 for the Kodak filter compared to 45.9 for the median filter and 14.6 for the MMSE filter). The slight increase in the chrominance MSE_{ab} values for the Kodak filter over the MMSE filter indicate that the Kodak filtered chrominance channels were not properly filtered as apparent from the color artifacts remaining in the flat regions of the filtered image.

To deal with these nonuniform color artifacts, which are mainly due to the severe noise artifacts that affect the chrominance channels [Figs. 9(a) and 9(b)], we use our

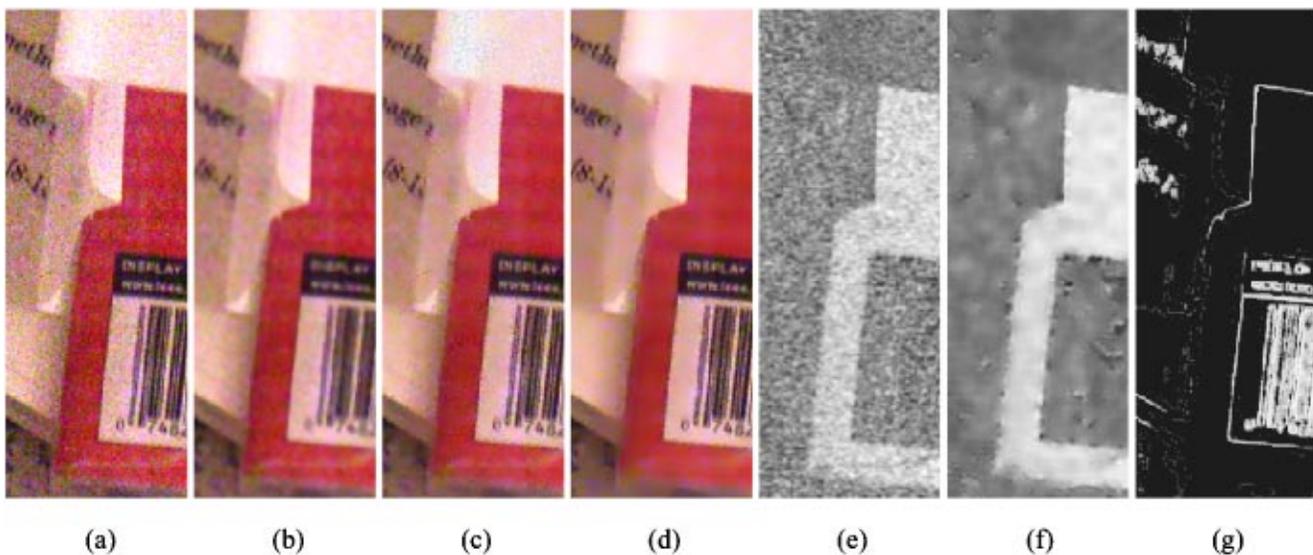


Fig. 10 (a) ISO-400 noise corrupted portion of the image in Fig. 3, with an estimated (L, a, b) -channel noise variance of $\sigma_n^2 = (41.9, 131.7, 160.5)$. (b) Median filtered image with 5×5 mask. Severe blurring of edges degrades the appearance of the image. (c) MMSE filtered image with 9×9 kernel mask. Color artifacts severely affect flat regions in the image. (d) HM-ACWMF filtered image. Color artifacts in flat regions have been reduced to a large extent because of proper filtering in the chrominance channels, while preserving edges. (e) Chrominance- a color channel with severe chromatic noise affecting the channel image. (f) HM-ACWMF filtered chrominance- a color channel, showing sufficient smoothing of color channel noise. (g) The values of the signal equalization estimator $\alpha(k, l)$.

HM-ACWMF with different filter parameters for the luminance and chrominance channel in the $L^*a^*b^*$ space. We used $\zeta=1$ with a maximum adaptive window size of 5×5 for the luminance L channel, and $\zeta=4$ with a maximum adaptive window size of 13×13 for the (a,b) channels with the idea of allowing more smoothing of the non-uniform color artifact noise in the chrominance (a,b) channels, while preserving details in the luminance L channel. This strategy worked well as the HM-ACWMF filtered image in Fig. 8(f) shows. The MSE_{Lab} has dropped to (8.7,14.6,37.3), lower than all the other filters, which shows that it gives the closest estimate to the reference image, both in terms of visual improvement, as well as noise reduction.

Last, but not least, we present results from filtering the noisy image of Fig. 3, which was acquired in a poorly lit environment with the digital camera CCD sensitivity set to ISO-400. Figure 10(a) shows a portion of the ISO-400 noisy image, which is severely corrupted with high-ISO noise of estimated (L,a,b) -channel noise variance $\sigma_n^2 = (41.9,131.7,160.5)$. Figure 10(b) shows a 5×5 median filtered image with severe smoothing of edges. Figure 10(c) shows a 9×9 MMSE filtered image with insufficient chromatic noise removal as apparent by the color artifacts in flat regions. Figure 10(d) shows our HM-ACWMF filtered image with a luminance maximum window size of 13×13 , and $\zeta=0.1$, and a chrominance maximum window size of 25×25 and $\zeta=0.4$. We used the same minimum window size of 5×5 for all channels. Figure 10(e) shows the noisy chrominance- a channel with severe chromatic ISO noise. Figure 10(f) shows the HM-ACWMF filtered chrominance- a channel showing significant chromatic noise removal. Figure 10(g) is an image of the values of the signal equalization estimator $\alpha(k,l)$. Bright areas of the image correspond to large values of α indicating large signal activity such as edges. Black areas are values of α near zero, indicating minimum signal activity. The α variable is an accurate edge detector as shown in the image.

6 Dealing with Long-Exposure Stuck-Pixel Noise

The earlier formulated HM-ACWMF together with the optimal-size windowing framework was shown to be very effective at reducing, and in less severe cases, completely removing, the last two types of high-ISO sensor noise stated in Sec. 2 earlier, namely blue-channel sensor noise (severely affecting the chrominance channels) and JPEG blocking artifacts.

In this section, an extra prefiltering step is introduced in our sensor noise filtering pipeline to deal with the stuck-pixel-type noise generated due to exposing the digital camera imaging sensor to the scene for an extended period of time [usually in dim light conditions such as for night shots as in Fig. 2(a)].

Methods of identifying noise pixels by using uncertainty measures have been reported in the literature.^{35,36} A method for noise filtering using contrast entropy was reported by Beghdadi and Khellaf.³⁷ We follow the idea of Beghdadi and Khellaf, but formulate the probability of a stuck-pixel using a local variance measure instead. This significantly reduces the blurring effects compared to the lower-order

contrast measure used in Ref. 37. We define the probability that pixel $y(k,l)$ centered in window $W(k,l)$ is a stuck pixel as

$$P(k,l) = \frac{V(k,l)}{\sum_{i,j \in W} V(i,j)}, \quad (13)$$

where

$$V(k,l) = |y(k,l) - \bar{y}|^2 \quad (14)$$

is a local gradient variance measure and \bar{y} is the mean pixel value inside the $N \times N$ window $W(k,l)$. The probability that this pixel $y(k,l)$ is a stuck pixel is, thus, $V(k,l)$ divided by the sum of variances of all pixels inside the window $W(k,l)$ as given in Eq. (13). Assuming that all pixels $\in W(k,l)$ are equally likely to be stuck pixels, then the probability of the window pixels being stuck-pixels can be given by $P_w = 1/N^2$, where the denominator denotes the total number of pixels in the window $W(k,l)$. This probability corresponds to a window where all the local gradient variances are equally distributed.³⁷ The criteria for a stuck pixel, thus, reduces to testing for a probability that is greater than P_w .

In our filter implementation, a new long-exposure flag is added to the algorithm, which, when set by the user, will indicate that the camera is in long-exposure mode and that the stuck-pixel prefilter (SPPF) is to be applied. Therefore, the probability of a stuck pixel is estimated for each input pixel $y(k,l)$ of the acquired image and, if $P(k,l) > P_w$, the stuck pixel $y(k,l)$ is removed by assigning to it the median value of pixels in a fixed 3×3 window centered around $y(k,l)$ and having values different from $y(k,l)$. On the other hand, if $P(k,l) < P_w$, indicating that the input pixel $y(k,l)$ is not a stuck pixel, then the HM-ACWMF is applied to this pixel for usual ISO-noise removal.

Figure 11(a) is the same simulated color edge image of Fig. 6(a). Here we use it again to test the performance of the SPPF by adding simulated stuck-pixel noise in the form of superimposed impulse noise on top of the Gaussian noise, as shown in Fig. 11(b). The total added MSE_{Lab} noise variance was calculated at (50.3,128.9,176.9). The median filtered image is shown in Fig. 11(c), with a $MSE_{Lab} = (3.1,14.9,16.5)$. Figure 11(d) gives the MMSE estimated image, which shows a complete failure to remove the simulated stuck pixels, which are regarded as edge features by the MMSE filter. The MMSE filter MSE_{Lab} values of (16.5,17.1,19.3) are higher than the median filtered values. Figure 11(e) shows the SPPF applied to the noisy image as a prefilter to remove the simulated stuck pixels, after which the HM-ACWMF is applied to remove the simulated high-ISO noise. Simple visual inspection reveals sharp edges with efficient noise removal for an overall superior visual quality compared to the median and the MMSE estimates. Also, the MSE_{Lab} values are the lowest at (0.6,11.2,10.3).

Figure 12(a) is an enlarged portion of the long-exposure image of Fig. 2(a) showing severe stuck-pixel noise due to long exposure at night. Figure 12(b) shows the image filtered using a fixed 3×3 window median filter. Excessive

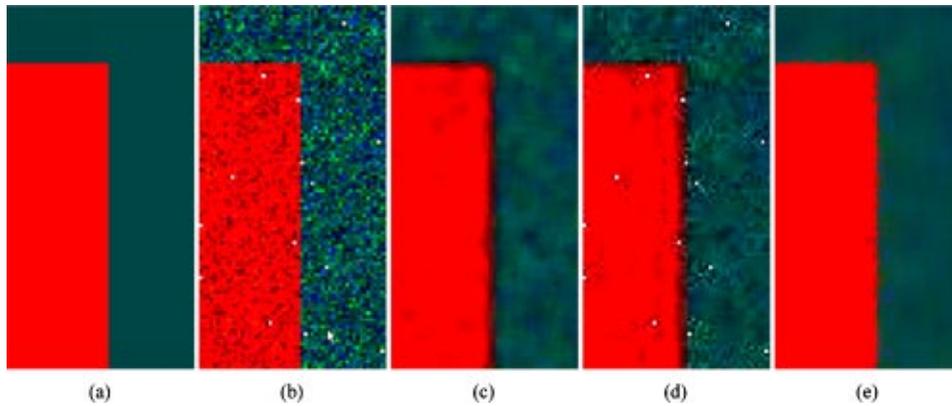


Fig. 11 (a) Ideal simulated color edge image. (b) Simulated long exposure stuck pixel noise plus simulated ISO noise, $MSE_{Lab}=(50.3,128.9,176.9)$ between the ideal and this noisy image. (c) 5×5 median filter, $MSE_{Lab}=(3.1,14.9,16.5)$. (d) MMSE filtered image (5×5 kernel mask), $MSE_{Lab}=(16.5,17.1,19.3)$. (e) SPPF followed by HM-ACWMF with a maximum adaptive window size of 11×11 , and $\zeta=0.1$. $MSE_{Lab}=(0.6,11.2,10.3)$.

smoothing of edges and loss of fine detail is apparent in the filtered image. Figure 12(c) shows the image filtered using our SPPF before applying the HM-ACWMF high-ISO noise filter. The SPPF window size used was 5×5 for testing the probability of a stuck-pixel, and 3×3 for computing and assigning the median to the stuck pixel. The adaptive window size used for the HM-ACWMF has an upper limit of 7×7 and a lower limit of 1×1 . It is clear that the stuck-pixel effects have been reduced to a large extent while preserving as much fine detail as possible compared to the median filtered image as apparent from the sharper visual results, which is an important concern in digital photography.

7 Conclusions

This paper has presented a hybrid noise removal filter specifically suited to the types of noise generated by digital photography. The issue of high-ISO noise in digital cameras was discussed, and three types of digital camera noise were identified. The hybrid mean and adaptive center weighted median filter was derived and an optimal-size windowing framework was implemented and used with the noise removal filter to reduce the effects of blue-channel noise and JPEG blocking artifacts common in high-ISO digital camera images. The third type of camera noise, which affects long-exposure images and causes stuck-pixel



Fig. 12 (a) Long exposure stuck pixel noise (courtesy <http://www.dpreview.com/>), (b) 3×3 median filter, and (c) SPPF (maximum adaptive window size used 7×7).

noise, was dealt with by preprocessing the image with a new method based on estimating the probability of a stuck pixel using local variance-based measures. This stuck-pixel filter was shown to be highly effective as a prefilter for removing stuck-pixel noise from long-exposure images. One observation from the experimental results is the fact that different sensors exhibit different noise levels at the same ISO setting. Finally, the methods developed in this paper were an attempt at addressing the growing concern in the digital photography community about the reduced visual fidelity in images acquired by modern professional digital cameras at high-ISO and long-exposure settings.

Acknowledgments

The author would like to thank the reviewers for their useful comments that helped improve the quality of the work described herein. This research work was financially supported by the Research Affairs at the UAE University, under contract number 01-01-9-11/04.

References

1. E. F. Haugh, "A structural theory for the Selwyn Granularity Coefficient," *J. Biol. Photogr. Assoc.* **11**(65), (1963).
2. G. C. Higgins and K. F. Stultz, "Experimental study of the RMS granularity as a function of scanning spot size," *J. Opt. Soc. Am.* **49**(925), (1959).
3. F. Naderi, "Estimation and detection of images degraded by film-grain noise," Tech. Rep. USCIP1 690, Univ. of Southern California, Los Angeles, California, Sept. 1976.
4. K. R. Castleman, *Digital Image Processing*, Prentice Hall, Upper Saddle River, NJ, 1996.
5. A. D. Hillery and R. T. Chin, "Iterative wiener filters for image restoration," *IEEE Trans. Signal Process.* **39**(8), 1892–1899 (1991).
6. M. Sonka, V. Hlavac, and R. Boyle, *Image Processing, Analysis, and Machine Vision*, 2nd ed., PWS, 1999.
7. J. M. Park, W. J. Song, and W. A. Pearlman, "Speckle filtering of SAR images based on adaptive windowing," *IEE Proc. Vision Image Signal Process.* **146**(4), 191–197 (1999).
8. H. C. Andrews and B. R. Hunt, *Digital Image Restoration*, Prentice-Hall, Englewood Cliffs, NJ, 1977.
9. K. N. Plataniotis, D. Androustos, S. Vinayagamoorthy, and A. N. Venetsanopoulos, "Color image processing using adaptive multichannel filters," *IEEE Trans. Image Process.* **6**(7), 933–949 (1997).
10. J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli, "Image denoising using gaussian scale mixtures in the wavelet domain," Tech. Rep. TR2002-831, Courant Institute of Mathematical Sciences, New York University, September 29 2002.
11. B. Smolka, K. W. Wojciechowski, and M. Szczepanski, "Random-walk approach to image enhancement," in *10th International Conference on Image Analysis and Processing*, September 1999.
12. H.-L. Eng and K.-K. Ma, "Noise adaptive soft-switching median filter," *IEEE Trans. Image Process.* **10**(2), 242–251 (2001).
13. C. A. Pomalaza-Raez and C. D. McGillem, "An adaptive, nonlinear edge-preserving filter," *IEEE Trans. Acoust., Speech, Signal Process.* **32**(3), (1984).
14. H. Kong and L. Guan, "Enhancement and real-time analysis of an adaptive impulse noise removal method," in *Proceedings of the 2nd IEEE International Conference on Engineering of Complex Computer Systems (ICECCS '96)*, pp. 164–167, 1996.
15. A. Flaig, G. R. Arce, and K. E. Barner, "Affine order statistic filters: A data-adaptive filtering framework for nonstationary signals," in *Proceedings of the 1997 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP '97)*, pp. 2145–2148, 1997.
16. S. E. Umbaugh, *Computer Vision and Image Processing: A Practical Approach Using CVIPtools*, Prentice Hall, Upper Saddle River, NJ, 1998.
17. K. Egiazarian, V. Katkovnik, and J. Astola, "Adaptive window size image denoising based on ICI rule," in *Proc. ICASSP'01*, pp. 1869–1872, May 2001.
18. <http://www.kodak.com>.
19. For detailed information about the CIE color spaces please visit their website at <http://members.eunet.at/cie/>.
20. The formulas for converting between color spaces are out of the scope of this paper, but can be obtained from: <http://www.easycrgb.com/math.php>.

21. Apple Computer Inc., *Advanced Color Imaging on the Mac OS*, Apple Computer Inc., November 1996, <http://developer.apple.com/techpubs/mac/ACI/ACI-2.html>.
22. A. B. Watson, "Perceptual-components architecture for digital video," *J. Opt. Soc. Am. A* **7**(10), 1943–1954 (1990).
23. S. Mann, "Comparametric equations with practical applications in quantigraphic image processing," *IEEE Trans. Image Process.* **9**(8), 1389–1406 (2000).
24. T. S. Huang, G. J. Yang, and G. Y. Tang, "A fast two-dimensional median filtering algorithm," *IEEE Trans. Acoust., Speech, Signal Process.* **27**(1), 13–18 (1979).
25. J. T. Astola and T. G. Campbell, "On computation of the running Median," *IEEE Trans. Acoust., Speech, Signal Process.* **37**(4), 572–574 (1989).
26. L. Khriji and M. Gabbouj, "Vector median-rational hybrid filters for multichannel image processing," *IEEE Signal Process. Lett.* **6**(7), 186–190 (1999).
27. M. Meguro and A. Taguchi, "The learning type of mean and median hybrid filters," in *Proceedings of the 1997 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP '97)*, pp. 2589–2592, 1997.
28. N. Himayat and S. A. Kassam, "A structure for adaptive order statistics filtering," *IEEE Trans. Image Process.* **3**(3), 265–280 (1994).
29. A. Taguchi and Y. Murata, "The median and the mean hybrid filters," in *Proc. IEEE Conference on Circuits and Systems*, pp. 93–96, 1991.
30. S.-J. Ko and Y. H. Lee, "Center weighted median filters and their applications to image enhancement," *IEEE Trans. Circuits Syst.* **38**(9), 984–993 (1991).
31. T. Chen and H. R. Wu, "Application of partition-based median type filters for suppressing noise in images," *IEEE Trans. Image Process.* **10**(6), 829–836 (2001).
32. P. Chan and J. S. Lim, "One-dimensional processing for adaptive image restoration," *IEEE Trans. Acoust., Speech, Signal Process.* **33**(1), 117–129 (1985).
33. *Applications of Digital Signal Processing*, A. V. Oppenheim, Ed., Chap. 4, pp. 169–237, Prentice-Hall, Englewood Cliffs, NJ, 1978.
34. Y. Boykov, O. Veksler, and R. Zabih, "A variable window approach to early vision," *IEEE Trans. Pattern Anal. Mach. Intell.* **20**(12), 1283–1294 (1998).
35. B. R. Frieden, "Restoring with maximum likelihood and maximum entropy," *J. Opt. Soc. Am.* **62**, 511–518 (1972).
36. A. Khellaf, A. Beghdadi, and H. Dupuisot, "Entropic contrast enhancement," *IEEE Trans. Med. Imaging* **10**(4), 589–592 (1991).
37. A. Beghdadi and A. Khellaf, "A noise-filtering method using a local information measure," *IEEE Trans. Image Process.* **6**(6), 879–882 (1997).



Tamer Rabie received his BSc honors degree in electronics and communications engineering in 1989 from Cairo University, his MSc degree in adaptive image processing in 1993 from the Department of Electrical and Computer Engineering at The University of Calgary, Canada, and his PhD degree in animat vision in January 1999 from the Department of Electrical and Computer Engineering at the University of Toronto. From October 1998 to December 1999 he held a postdoctoral fellow in the Department of Computer Science at the University of Toronto where he was involved in research and development of the animat vision project. From January 2000 to August 2001 he held an assistant professor appointment in the Department of Electrical and Computer Engineering at Ryerson University in Toronto, Canada. He is presently an assistant professor in the College of Information Technology at the United Arab Emirates University since September 2001. He is also an adjunct professor in the Department of Civil Engineering at the University of Toronto's Intelligent Transportation Systems Centre, conducting collaborative research work in active computer vision-based traffic surveillance and control since May 2000. His current research interests include intelligent transportation systems (ITS); virtual reality and its application to simulated computer vision; digital image processing; and computer graphics and visualization. Dr. Rabie has published papers in ITS, image processing, computer vision, and artificial intelligence. Dr. Rabie is a senior member of the IEEE, and has an active membership in the IEEE Signal Processing Society, the IEEE Computer Society, and the Professional Engineers of Ontario (PEO) Association in Canada.