

CCD Noise Removal in Digital Images

Hilda Faraji, *Member, IEEE*, and W. James MacLean, *Senior Member, IEEE*

Abstract—In this work, we propose a denoising scheme to restore images degraded by CCD noise. The CCD noise model, measured in the space of incident light values (light space), is a combination of signal-independent and signal-dependent noise terms. This model becomes more complex in image brightness space (normal camera output) due to the nonlinearity of the camera response function that transforms incoming data from light space to image space. We develop two adaptive restoration techniques, both accounting for this nonlinearity. One operates in light space, where the relationship between the incident light and light space values is linear, while the second method uses the transformed noise model to operate in image space. Both techniques apply multiple adaptive filters and merge their outputs to give the final restored image. Experimental results suggest that light space denoising is more efficient, since it enables the design of a simpler filter implementation. Results are given for real images with synthetic noise added, and for images with real noise.

Index Terms—Camera response function (CRF), light space, nonlinear filtering and enhancement, noise modeling, restoration.

I. INTRODUCTION

TYPICALLY, CCD imaging incorporates three main components: 1) an image sensor array, which converts light into voltage at pixels, 2) camera electronics, which impose a nonlinear function on the voltage values and act as a dynamic range compressor, and 3) an analog to digital conversion unit, which generates the final image values. These steps are shown in Fig. 1. Light photons arrive at a CCD sensor cell and are converted to voltage. While, ideally, this voltage would be linearly related to the quantity of photons, in reality, nonideal CCD behaviour and subsequent processing by the camera (for example application of nonlinear transfer functions) mean the mapping between incident photons and camera output is nonlinear, and is given by a function called the *camera response function* (CRF). Finally, the voltage values are quantized by an A/D conversion unit.

However, as shown in Fig. 1, CCD sensor output not only carries the useful signal, but it also includes a variety of noise components which can seriously restrict the ability to achieve high-quality images. The objective of this research is to design a postprocessing software-based algorithm to restore images contaminated by the CCD noise sources.

A. Light Space Versus Image Space

CCD sensor output values, before processing and digitization, lie in a linear space called “light space” (LS) [1]. Light space values are denoted as q , as shown in Fig. 1. Image space (IS)

Manuscript received June 8, 2004; revised January 2, 2006. This work was supported by NSERC. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Nicolas Rougon.

The authors are with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON M5S 3G4, Canada (e-mail: hilda@eecg.toronto.edu; maclean@eecg.toronto.edu).

Digital Object Identifier 10.1109/TIP.2006.877363

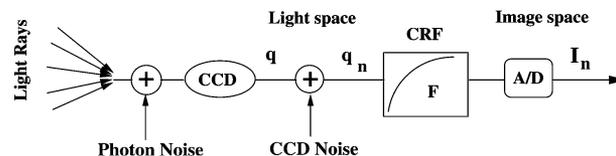


Fig. 1. Basic components of a typical CCD imaging system. CRF represents the camera response function.

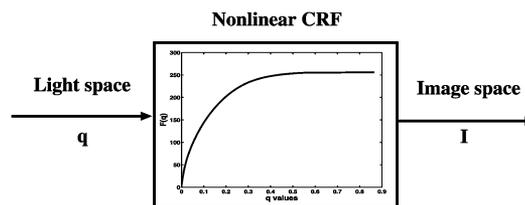


Fig. 2. CRF performs a nonlinear mapping between light space and image space.

refers to the final pixel values, I , which are nonlinearly related to light space values, q , for most digital cameras.

The CRF performs the nonlinear mapping¹ between the two spaces such that $I = \mathcal{F}(q)$ and $q = \mathcal{F}^{-1}(I)$ as shown in Fig. 2. The CRF is shown by plotting q against image intensity I . For very small q values, the function is almost linear with steep slope. As q increases, the nonlinearity becomes more evident and it levels off or saturates for high q values (Fig. 2).

The ground-truth CRF can be accurately obtained by using a standard test pattern [2]. Multiple images of the test chart with different exposure times can be used to obtain multiple pairs of $(q, \mathcal{F}(q))$. Since the reflectance of each portion is known, q can be assumed to be a known parameter. Also, $\mathcal{F}(q)$ is simply the corresponding pixel value. Corresponding to any exposure time, a CRF can be estimated. The final step to estimate a CRF involves the proper alignment of these response functions [2]. Throughout this work, we have used the Kodak DCS260 ground-truth CRF for the experiments, except for the presentation of the images with real noise captured from a Nikon D70.

The response function can also be estimated even if the standard test chart is not available. In such nonparametric methods, multiple images of the same subject matter with different exposure values (k_i) are collected to construct a “comparagraph” [2], in which $\mathcal{F}(k_i q)$ is plotted against $\mathcal{F}(q)$. The response function is recovered using an unrolling process [1]. A small value, q_0 , is initially considered as a reference for which $\mathcal{F}(q_0) = q_0$. The corresponding value on the vertical axis results in $\mathcal{F}(k q_0)$. If this value is found along the x axis, the comparagraph gives $\mathcal{F}(k \cdot k \cdot q_0) = \mathcal{F}(k^2 q_0)$. Continuing this process determines multiple pairs of the form $(q, \mathcal{F}(q))$, where q refers to a set of

¹While most CRFs are nonlinear, some cameras offer a linear CRF through a user-selectable “gamma function.”

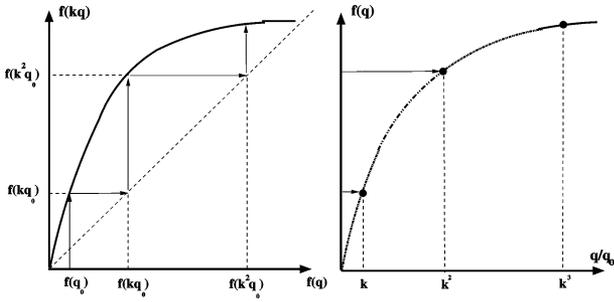


Fig. 3. Plot of $f(kq)$ versus $f(q)$ and (left) the unrolling process. (Right) Final response function [2]. Parameter f represents the CRF.

known values. The final CRF is in the form of a lookup table, which can be used to convert images into light space, and vice versa. The CRF measured for a given camera may be somewhat inaccurate if the unrolling estimation technique is used, as noise can cause the comparagraph to “thicken” to a certain extent. Estimation from a standard test pattern is more accurate. This process, called “unrolling,” is shown in Fig. 3.

B. CCD Noise

As shown in Fig. 1, CCD sensor output not only carries the useful signal but it also includes a variety of noise components such as photon noise, fixed-pattern noise (FPN), and amplifier noise [3]. Photon noise is related to the random fluctuation of photon flux arriving at the CCD sensor, and it follows Poisson statistics. FPN is due to differences in individual pixels’ responsivities. This type of noise is more prominent at higher intensities and is signal dependent. Its statistics are proportional to the original signal. Readout noise is introduced to the signal during the process of measuring the signal. This type of noise is modeled as white noise.

Typically, restoration algorithms assume a linear mapping between the incident LS and IS [4]. In practice, the nonlinear CRF causes the sensor noise model to become more complex in IS. Despite the importance of nonlinear CRFs on image restoration, its application has not been stressed in the literature. In Pavlovic and Tekalp [5], a Wiener filter is applied for the restoration of images recorded by photographic film. In their proposed method the nonlinearity between the incident light on the film and the recorded images is incorporated into their restoration algorithm. The noisy images are transformed into “exposure domain” in which there is a linear relationship between the image intensity and the incident light. Exposure domain is similar to light space in that the sensor output is linearly dependent on the incident light.

A linear physical model for a CCD sensor is used to quantify the CCD noise in [6]. Healey *et al.* assume a linear mapping between the space of sensor output and image space. The total noise is considered as a combination of dark current noise, shot noise, readout noise, and quantization noise. They then subdivide noise into signal-dependent and signal-independent components. Each noise component is then estimated based on the prior knowledge of its distribution. Tsin *et al.* [7] modify the model presented by Healey [6] by incorporating a nonlinear function in the estimation process. The developed model is used

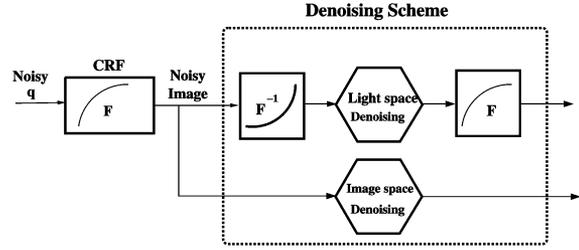


Fig. 4. Our two denoising schemes.

to estimate the CRF. The estimate of this function is through a maximum-likelihood procedure in an iterative fashion. The CCD output is then modeled as a combination of noise-free sensor output, shot noise, and thermal noise.

In this work, our restoration algorithm is based on the actual CCD noise model by taking advantage of LS in which the q values represent the sensor output, and, therefore, the LS noise model resembles the actual noise model. Further details are given in Section II.

The two main sections of this paper involve 1) finding the CCD noise model and 2) designing a denoising algorithm based on the estimated CCD noise model. This paper begins with a general description of our filtering approach, followed by simulation results in which we compare the performance of the two methods.

II. OUR GENERAL RESTORATION APPROACH

Our general denoising scheme is shown in Fig. 4. The approach is summarized as 1) off-line modeling of image and CCD noise, and 2) developing two separate restoration methods based on the noise model in LS and on the nonlinear transformed version of this model in IS. We finally evaluate the performance of the filters and compare the results in IS and LS.

Filtering in LS involves three main stages [8]. A noisy image is initially converted to LS by applying the inverse of the CRF, \mathcal{F}^{-1} . In LS, the image is processed using our denoising technique. The filtered image is then converted back to IS. In the second approach, the noise model in IS is estimated by transforming the LS noise model to IS. The noisy images are then processed based on noise characteristics in IS. Finally, the performance of these two methods are compared in order to determine any benefits of LS denoising.

A. CCD Noise Model

In this section, we will explain the estimation of the CCD noise model in both LS and IS. For CCD systems, a photon transfer curve (PTC) is a calibration standard used to electronically model CCD noise [3]. A PTC is a plot of noise standard deviation against the CCD sensor output (electrical charge), at different illuminations, on a logarithmic scale. Estimating the PTC typically requires hardware experiments. However, LS approximates the sensor output; thus, the LS noise curve resembles a typical PTC as shown in Fig. 5. The PTC is mainly organized into three categories. At the lowest signal level, the curve is almost flat, indicating that noise for low intensities is independent of the signal. Amplifier noise is the dominant noise in this region [3]. Increasing the illumination causes noise to become signal

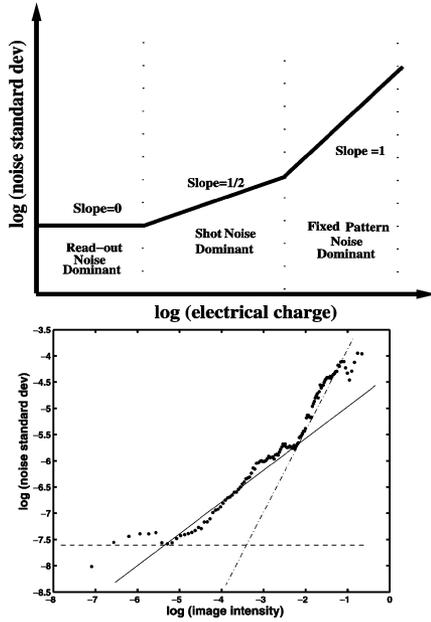


Fig. 5. (Top) An ideal PTC. (Bottom) LS noise curve on the logarithmic scale for a Kodak DCS260 camera (experimentally determined in the authors' lab for a specific camera). Dotted lines denote the linear approximations to the three regions. The q values are normalized to the range $[0, 1]$.

dependent. A mild upward trend for middle intensity range has the slope of approximately one half. The dominant noise in this region is shot noise. This type of noise is mainly a combination of photon noise and dark noise, both exhibiting Poisson distributions since it is based on the random arrival of photons at the CCD [3], although this distribution can be approximated as Gaussian [9]. The last region, corresponding to high image intensities, is dominated by FPN. This type of noise is proportional to the signal, and, therefore, this region exhibits a slope of unity. Since FPN can be removed by a preprocessing operation called “flat fielding” [3], we can reasonably assume a two-region PTC. Different cameras are expected to generate similar PTCs, i.e., three distinct noise regions as mentioned above [3].

In IS the noise model is estimated by applying the nonlinear CRF to the test images used to estimate the PTC in LS. We then compute IS noise standard deviation versus noise-free image values. Fig. 6 shows an example of a simulated noise curve in LS and its nonlinear version in IS. The simple two-region noise model is transformed to a more complex noise model, which includes at least three noise regions. Generally, noise characteristics are different in IS, with the derivative of the CRF exerting the main influence. Let n_q and n_I denote the noise in LS and IS respectively. We know

$$n_I = I_n - I = \mathcal{F}(q + n_q) - \mathcal{F}(q). \quad (1)$$

Using the Taylor series expansion of \mathcal{F} around q results in

$$\begin{aligned} \mathcal{F}(q + n_q) &\approx \mathcal{F}(q) + \dot{\mathcal{F}}(q)n_q \\ \Rightarrow n_I &\approx \dot{\mathcal{F}}(q)n_q. \end{aligned} \quad (2)$$

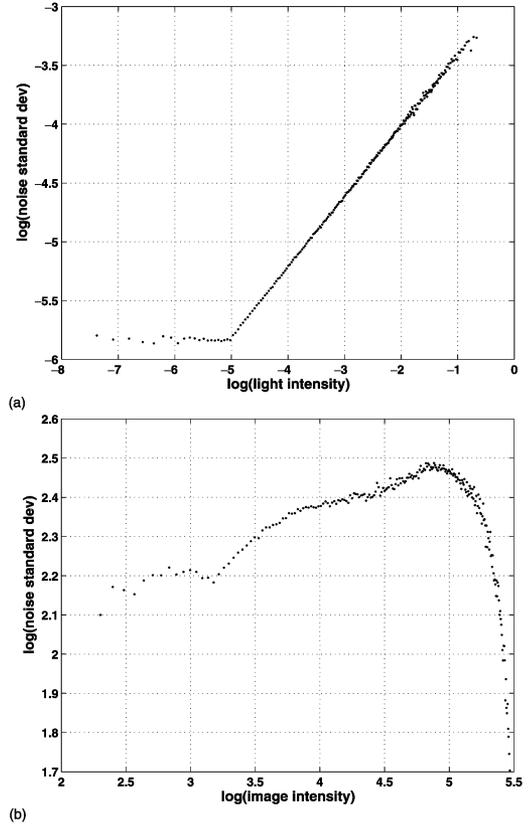


Fig. 6. (a) Noise model in LS and (b) noise model in IS.

For very low intensities, the derivative is almost constant and, therefore, noise behavior is similar in both domains. Reduction in $\dot{\mathcal{F}}$ causes a decrease in IS noise. A sharp decrease in the noise curve for high intensities is due to $\dot{\mathcal{F}}$ approaching zero in the saturation area. Here, we ignore quantization noise as, through a series of experiments, it appears that the quantization operation does not substantially affect noise distribution in the image space. As shown in Fig. 6, CCD noise variance appears to be a function of q and, therefore, signal-dependent in both image and light domains. To generate the noise function based on its variance, each noise region is approximated by a straight line with known offset σ_u with respect to the origin and known slope γ . As shown in Fig. 6 (top), noise in the second region, where intensities are higher than e^{-5} , is characterized as

$$\begin{aligned} \log \sigma_{n2} &= \gamma \log(q) + \log(\sigma_{u2}) \\ \rightarrow \sigma_{n2} &= \sigma_{u2} \cdot q^\gamma \end{aligned} \quad (3)$$

where σ_{n2} and σ_{u2} are noise standard deviation and offset for the second noise region. For instance, let σ_c denote the noise standard-deviation in the first region of the PTC. Ignoring FPN, the total noise standard deviation, σ_n , associated with a typical CCD sensor can be formulated as

$$\sigma_n = \begin{cases} \sigma_c & f \leq e^B \\ \sigma_u \cdot f^\gamma & f > e^B. \end{cases} \quad (4)$$

where e^B refers to the breakpoint between the two noise region [B is -5 in Fig. 6 (top)]. When the log-intensity is less than a given value, e^B , noise can be modeled as constant.

Most types of signal-dependent noise (SDN) can be expressed in a parametric form as

$$\eta(n) = \sigma_u \cdot f^\gamma(n) \mathcal{N}(n) \quad (5)$$

which is then added to the noise free signal as

$$g(n) = f(n) + \eta(n) \quad (6)$$

where the index n denotes the image coordinate, g refers to the noisy image, and f is the unknown noise-free image [9]. Sensor noise η is defined as a function of f and a zero-mean, unity-variance Gaussian noise \mathcal{N} . Noise parameters σ_u and γ are determined based on the noise curve as explained earlier. While experiments [3] performed directly on CCD devices confirm $\gamma = 0.5$, noise curves measured by the authors on entire camera systems had values that ranged from 0.2 to 0.5. Once the noise variance is known, the final noise function is modeled as denoted in (6) and (5).

B. Restoration Algorithm

Our denoising methods are based on linearized noise curves which include multiple noise regions. Instead of assuming a single noise model for all intensities, our technique involves applying multiple filters, each of which is optimal for a specific noise region, and an intelligent combination of the resulting outputs. We describe a general restoration algorithm derivation that is applicable in both LS and IS.

Many attempts have been made to restore images degraded by SDN [10]–[12]. In [11], an adaptive noise smoothing (ANS) filter is applied to images degraded by SDN. This local filter, operating on a predefined neighborhood, is based on the nonstationary mean, nonstationary variance (NMNV) image model. In [12], an ANS filter is applied in the wavelet domain to restore images degraded by SDN whose model is denoted in (5). Uniform patches are determined from the observed (noisy) image, and used to estimate the noise parameters γ and σ_u [13]. However, detecting uniform image regions in the noisy image is challenging when dealing with SDN and strong assumptions are made. Consequently, the estimated noise parameters may not be accurate. In this work, we accurately estimate the noise parameters using the observed noise curves obtained through an off-line method. We apply ANS filters as explained in [11] and [12]. The ANS filter formulation for SDN can be written as

$$\hat{f}_{\text{ANS}}(n) = \left(1 - \frac{\sigma_f^2(n)}{\sigma_g^2(n)}\right) \mu_f + \frac{\sigma_f^2(n)}{\sigma_g^2(n)} \cdot g(n) \quad (7)$$

where n refers to the pixel index, \hat{f}_{ANS} is the noise-free estimate, g is the noisy image, σ_g^2 is the noisy local variance, and μ_g is local mean. As indicated by (7), the filter adapts to local statistical variations such that uniform regions are smoothed, while edges and fine details are invariant to the averaging process.

When the local image region is uniform, the estimated $\sigma_f^2(n)$ is considerably smaller than the corresponding $\sigma_g^2(n)$, causing the ratio $\sigma_f^2(n)/\sigma_g^2(n)$ to be close to zero. Thus, the filter puts most of the weight on $\mu_g(n)$, smoothing the areas. Conversely, in the presence of a sharp edge, $\sigma_f^2(n)/\sigma_g^2(n)$ is close to one, and, as a result, the filter puts more weight on the noisy values, $g(n)$. In this case, the edge sharpness is preserved.

The noise-free local statistics, given by σ_f^2 and μ_f , need to be estimated based on the noisy local variance σ_g^2 , local mean μ_g , and prior knowledge of the noise model. It is important to note that the quality of the estimates for μ_g and σ_g^2 are strongly dependent on the size of the window used to estimate these parameters, and windows that are too small may lead to poor filter performance. The smallest window size we attempt is 3×3 . Since we assume that the Gaussian noise \mathcal{N} is independent of the noise-free image f , there is no correlation between f and the Gaussian noise; thus, $\mu_f = \mu_g$. Furthermore, $\sigma_g^2(n)$ can be written as

$$\sigma_g^2 = \sigma_f^2 + \sigma_u^2 E[f^{2\gamma}]. \quad (8)$$

We drop n for notational simplicity. To simplify the expression for $E[f^{2\gamma}]$, we use the Taylor series expansion of $f^{2\gamma}$ around μ_f , so

$$E[f^{2\gamma}] \approx \mu_f^{2\gamma} + \gamma(2\gamma - 1)\mu_f^{2\gamma-2}\sigma_f^2. \quad (9)$$

Finally, σ_g^2 can be summarized as

$$\begin{aligned} \sigma_g^2 &= \sigma_f^2 + \sigma_u^2 E[f^{2\gamma}] \\ &\approx \sigma_f^2 + \sigma_u^2 \left(\mu_f^{2\gamma} + \gamma(2\gamma - 1)\mu_f^{2\gamma-2}\sigma_f^2 \right) \\ \sigma_f^2 &\approx \frac{\sigma_g^2 - \sigma_u^2 \mu_f^{2\gamma}}{1 + \gamma(2\gamma - 1)\mu_f^{2\gamma-2}\sigma_u^2}. \end{aligned} \quad (10)$$

Therefore, \hat{f}_{ANS} is a function of σ_g^2 , μ_g , and the noise parameters (as also derived in [12])

$$\hat{f}_{\text{ANS}}(n) = \frac{\left(1 - \sigma_u^2 \frac{\mu_g^{2\gamma}(n)}{\sigma_g^2(n)}\right) (g(n) - \mu_g(n))}{1 + \gamma(2\gamma - 1)\mu_g(n)^{2\gamma-2}\sigma_u^2} + \mu_g(n). \quad (11)$$

1) *Edge Map*: In addition to the filter structure, the window size also largely affects the filter performance. As mentioned earlier, the ANS filter prevents edge blurriness by assigning less weight to μ_g in nonuniform regions; however, this creates artifacts in the vicinity of sharp edges. The distortion width around the edges depends on the window size of the filter. The bigger the window size, the broader the distorted region.

We deal with this problem by taking advantage of the fact that ANS filters with smaller window sizes create less artifacts, thus considerably reducing this effect. On the other hand, satisfactory performance in uniform areas improves significantly by increasing the window size. Therefore, we make window size an adaptive parameter so that the ANS filter operates on small

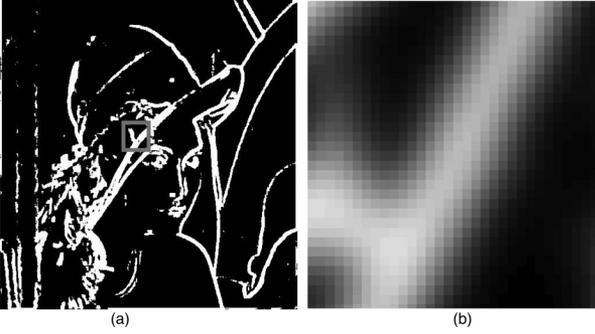


Fig. 7. (a) "Lena" edge map and (b) enlargement of the region indicated by the red box in (a). High intensities in edge map corresponds to the edges. Low intensities represent the uniform regions. The ANS filter with small size will operate on the high intensity regions.

neighborhoods in the vicinity of edges and increases its size for uniform areas. We then blend the outputs of these two filters by using a spatial edge map in the range of $[0, 1]$ in which pixels with intensities close to one represent the edge regions. Let \hat{f}_S and \hat{f}_L denote the output images corresponding to the two ANS filters with small and large window sizes. The nonlinear ANS filter f_{NANS} is then estimated as

$$\hat{f}_{\text{NANS}}(n) = EG(n)\hat{f}_S(n) + (1 - EG(n))\hat{f}_L(n) \quad (12)$$

where n represents the image coordinate, and EG is the edge map. Examples of edge maps are provided in [8] and Fig. 7. As indicated by (12), the results are combined so that the filter with smaller window size is dominant in edge regions, but gradually reduces its influence in uniform regions where the second filter becomes dominant.

The algorithm used to create the edge map is as follows; a binary edge map is initially generated by thresholding the ratio between the estimated σ_f^2 to σ_g^2 . This binary map is then smoothed by application of a grey-scale morphological dilation to let the edges diffuse throughout a small neighborhood, creating a smooth transition between the two ANS filters as shown in Fig. 7. The output image at this stage is the NANS filter.

2) *Intensity Map*: As shown in Fig. 5, CCD noise exhibits varying characteristics depending on image intensities. Therefore, we apply multiple NANS filters with different values of γ and σ_u on a single noisy image. For instance, to address the signal-independent noise, we set parameter γ equal to zero. At the final stage, we design an intensity map to smoothly merge the NANS outputs. The intensity map is composed of several weight functions, each of which is related to a single noise region. If there are K noise regions, we apply K NANS filters to the entire image. The weight functions W_k , which sum to one for any intensity, then assign different weights to different outputs, depending on the pixel intensity. The weight functions are based on the well-known sigmoid function, $f(x) = (1 + e^{-\alpha x})^{-1}$. For the example, shown in Fig. 8, the weight function dominant in the region $x \leq B_1$ is $w_1(x) = (1 + e^{\alpha(x-B_1)})^{-1}$, the weight function dominant in the region $x \geq B_2$ is $w_3(x) = 1 - (1 + e^{\alpha(x-B_2)})^{-1}$, and the weight function dominant in the region $B_1 < x < B_2$ is $w_2(x) = 1 - (w_1(x) + w_3(x))$. Here, α controls the steepness of the transition between regions, and for

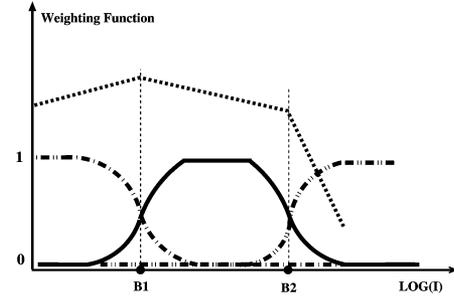


Fig. 8. (Dash-dot, solid, and dash-dot-dot curves) Three weighting functions provide a smooth combination of ANS filtered images around (dotted line) the noise curve knees e^{B_1} and e^{B_2} . Please note that the noise curve is included here in a purely schematic sense to explain the positioning of the breakpoints and is not related to the scale of the y axis. This weighting scheme can be used in light space by omitting the second breakpoint and reducing the number of weighting functions to two.

this example can be chosen as $\alpha \geq (B_2 - B_1)/2$. When well inside the m th noise region, W_m has a maximum value close to one, and other weight functions are close to zero. In the overlap regions, W_m monotonically reduces its value until the adjacent intensity functions become dominant [8], [14]. The final filtered image at the n th pixel is

$$\hat{f}_{\text{Final}}(n) = \sum_{k=1}^K W_k(n)\hat{f}_k(n) \quad (13)$$

where n is the image coordinate, and \hat{f}_k is the NANS filtered image corresponding to the k th noise region as indicated by (12). The technique discussed in this section can be applied in either LS or IS. An example of a weight intensity map is shown in Fig. 8.

III. FILTERING IN LIGHT SPACE

For applications where the raw camera response data is required, or where the defect in the image originates in LS, processing the sensor output directly leads to more satisfactory results [1], [15]. LS enables accessing the sensor output and correcting for nonlinearity in the IS noisemodel. This motivates applying our restoration algorithm in LS. Furthermore, the LS noise model resembles the standard electronic noise model, namely the PTC, which contains two noise regions, one signal-independent noise region and one SDN region. As shown in Fig. 9, a noisy image is initially converted to LS. It is then processed using a combination of two NANS filters and an adaptive Wiener filter, where the latter targets the dark regions in the image. An edge map is used to combine the results of ANS1 (small window) and ANS2 (big window). The filter parameters are determined by fitting the two-region noise model to the PTC as measured in light space. The final filter is the weighted average of individual output images using the developed intensity functions. In the last step the filtered image is transformed back into IS.

IV. FILTERING IN IMAGE SPACE

As shown in Fig. 6, the noise model in IS becomes more complex in that it usually contains three or more noise regions and

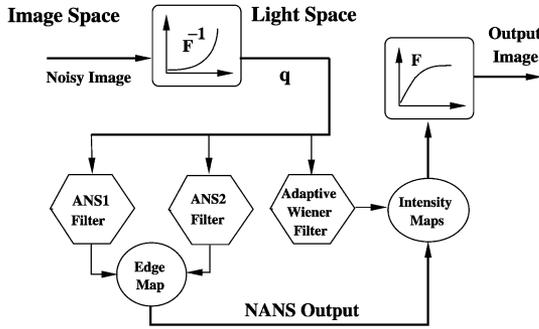


Fig. 9. Our general filtering scheme in light space. Corresponding to any region of the noise model, an adaptive smoothing filter is designed and applied to the entire noisy image. The final output image is a weighted average of the individual ANS outputs plus the Wiener filter output.

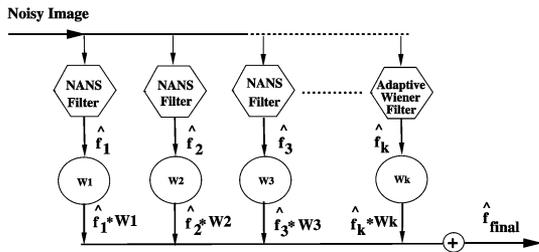


Fig. 10. Our general filtering scheme in image space. Corresponding to any portion of the noise model an adaptive smoothing filter is designed and applied to the entire noisy image. The final output image is a weighted average of the individual outputs.

it includes curve segments instead of straight segments. Therefore, the restoration is computationally more demanding than that in LS. We develop the IS denoising method in order to examine the benefits and drawbacks of each method and, therefore, enable performance comparisons.

The filter parameters for image space are determined by fitting the three-region (or multiregion) model to the PTC as measured in image space. When the noise curve is generated in IS, the breakpoints, partitioning different noise regions, are initially identified. Each curve segment is then regressed to a line segment. Noise parameters for each segment are found. Similar to LS, the final result is the combination of filtered images. The last noise region, having a large negative slope, is processed by an adaptive Wiener filter, since ANS filters do not handle this type of noise well. This happens since the last region in the image space noise curve (high intensities) results in a very large intercept, and it is not possible to apply ANS filters due to the numerical roundoff problems when σ_u is very large. It is important to note that in the presence of signal-dependent noise, Wiener filters do not achieve their optimal performance. Our general filtering system in image space is shown in Fig. 10.

V. EXPERIMENTAL RESULTS

We have performed experiments on a series of images corrupted by simulated noise based on the typical photon transfer curve. In this section we present results and discuss the performance of our denoising schemes in both image space and light space, followed by a performance comparison between the two methods. Finally, results from applying light-space denoising to a set of real images is given.

A. Real Images With Synthetic Noise

In order to estimate the PTC, we used a real scene and collected a large number of images at a fixed exposure level, using a Kodak DCS260 camera. The images were then transformed to light space and the average of the resulting images were computed in the temporal direction. We used the (known) ground-truth response function for this camera as estimated using a standard test pattern. The total noise variance was then computed as the variance of all samples corresponding to each intensity in the average image over all the frames. To construct the noise model in IS, we used the same procedure without any transformation to LS. The resulting plot is shown in Fig. 6. It is of interest to consider how much variability in the CRF would result in a family of cameras of the same model; however, a study of this is beyond the scope of the current paper. It is worth noting—although, again beyond the scope of this paper—that good denoising results are obtained for a range of filter parameter values σ_u and γ , so the method is not sensitive to small variations in these parameters.

Our restoration approach in LS is based on Fig. 9. For experiments, a set of noise-free images (standard images or an average of a large number of real images corresponding to the same scene) are initially converted to light space using the Kodak DCS260 response function. Synthetic noise, based on the model shown in Fig. 6, is then added to the images in light space. The PTC for both image space and light space are computed, and the filter parameters γ and σ_u are estimated from the curves (note: in order to compute the PTC in these experiments, we assume knowledge of the noise-free images). In the light-space case, γ and σ_u need only to be estimated for the region above B1. For the image-space case, separate γ and σ_u values are computed for the region between B1 and B2 and the region above B2. The constant noise region (below B1) just has parameter σ . For both light space and image space PTCs, B1 is automatically determined by the software. The value of B2, and γ and σ_u for the third region of the image space PTC, are manually determined.

We then process the degraded images by applying two ANS filters (with window sizes 3×3 and 9×9) and a standard adaptive Wiener filter in LS. The results are then combined using the weight intensity maps. As shown in Fig. 11, the performance of the ANS filter with small window in uniform regions is poor, but the edges are preserved. In contrast, the ANS with large window significantly smooths the uniform areas, but generates artifacts in edge regions. Combining the two ANS filters deals with the artifact problem.

Fig. 12 shows how the weight function combines the results of the NANS, and adaptive Wiener filters. As shown, the adaptive Wiener filter blurs the image, while the NANS filter performs poorly in dark regions. These problems are removed by applying this map on both output images and combining the results. Overall, these figures illustrate that our final combined filter outperforms adaptive Wiener results by removing its blurriness effect, ANS filter by reducing the artifacts around the edges, and NANS filter by reducing the remaining errors in the dark regions of an image.

We also compute mean square error (MSE), minimum absolute error (MAE), and signal-to-noise ratio (SNR) as quantitative noise-attenuation measures. SNR is computed as

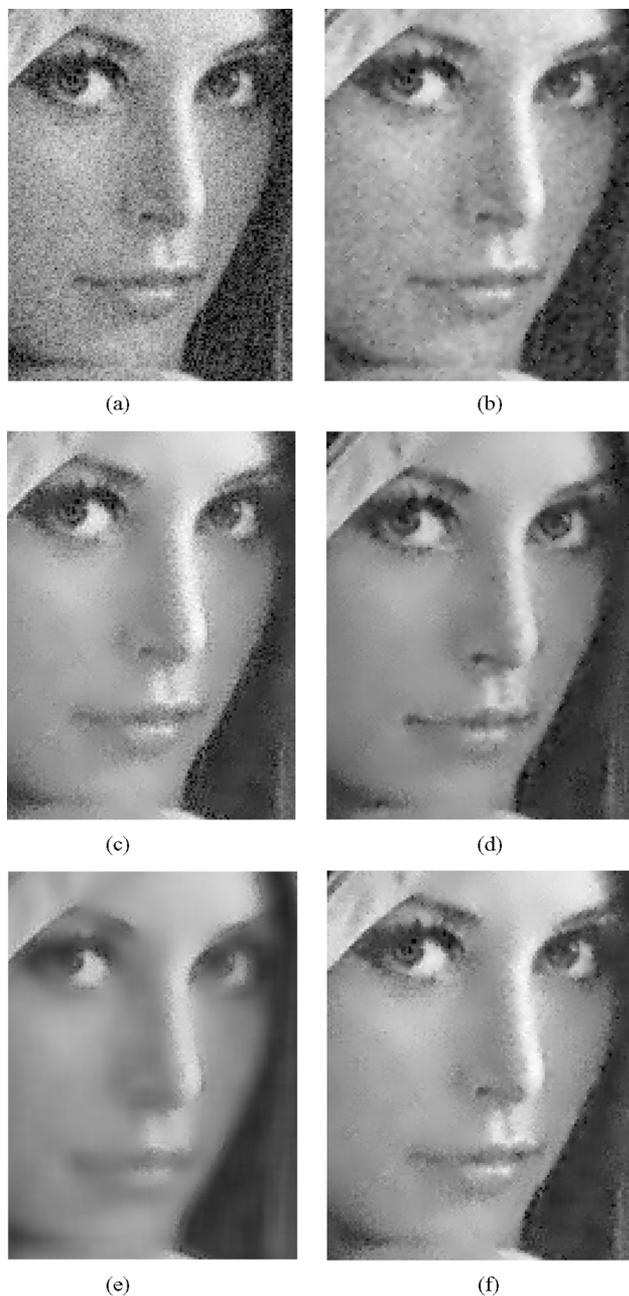


Fig. 11. Effect of edge map on artifact reduction: (a) corrupted “Lena” image with $\gamma = 0.3$ and $\sigma_u = 0.05$; (b) filtered image with a 3×3 ANS filter; (c) 9×9 ANS filtered image; (d) the NANS result; (e) filtered image with Wiener; (f) our final result. Since the image consists mainly of intensities higher than 50, the final image is similar to the result of ANS filter.

$10 \log_{10} \left\{ \frac{\sum_x f^2(x)}{\sum_x (f(x) - \hat{f}(x))^2} \right\}$ where \hat{f} is the filtered image, f is the noise-free image, and x are the pixel indices. Table I provides MSE values for various values of γ and compares the performance of adaptive Wiener, NANS, and the final combined filters. As shown, the noise level in the final combined image is significantly lower than the other filtered images.

To denoise in IS, as shown in Fig. 10, we convert the noisy image back to IS and apply the method of Section IV. Both

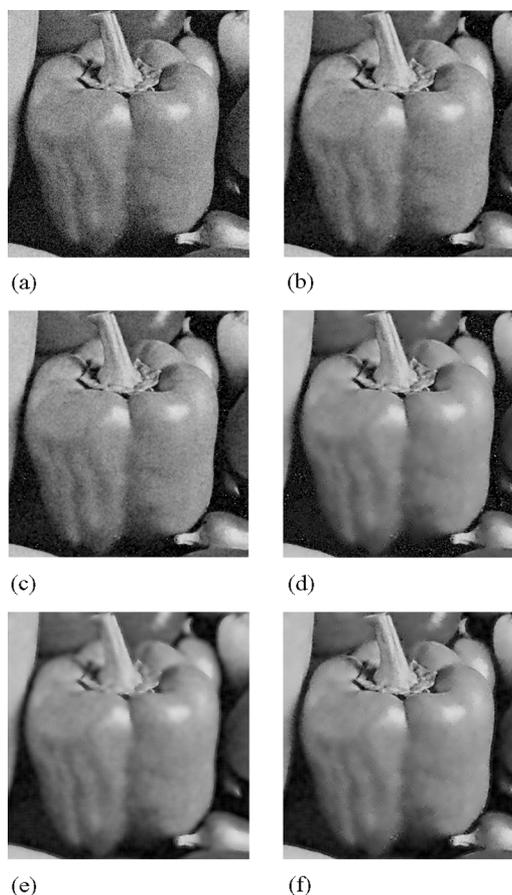


Fig. 12. Effect of the weight function: (a) cropped “Peppers” image degraded by $\sigma_u = 0.04$ and $\gamma = 0.3$; (b) filtered image with a 3×3 ANS filter; (c) 9×9 ANS filtered image; (d) NANS estimate; (e) adaptive Wiener estimate; (f) the final result, in which artifacts in dark regions are removed.

TABLE I
MSE MEASURES OF THE “PEPPERS” IMAGE FOR THREE VALUES OF γ , $\sigma_u = 0.04$. RESULTS FOR LS FILTERING

γ	Noisy	Wiener	NANS	Final combined
0.3	323.481	98.736	100.066	65.134
0.5	104.889	94.897	39.488	38.226
0.7	35.048	94.161	24.888	18.225

quantitative and qualitative measures are used to judge the performance of both algorithms. Uniform region and edge preservation are assessed visually.

Fig. 13 shows an example of noise-free and noisy images, together with the filtered images using a Wiener filter, a median filter and our algorithms. As shown, LS and IS filtering outperform the other filters in preserving fine details and sharp edges. It also appears that, except for bright areas, restoration in IS and LS visually yields equally satisfactory results in terms of noise smoothing and edge preservation. In bright areas LS filtering outperforms IS filtering, as seen in the bottom two images in Fig. 13. Table II provides a quantitative performance comparison of the adaptive Wiener filter and our methods for the “Lena” image. Further examples are available in [8].

Experimental results suggest that denoising in LS usually results in slightly higher quantitative performance over the whole

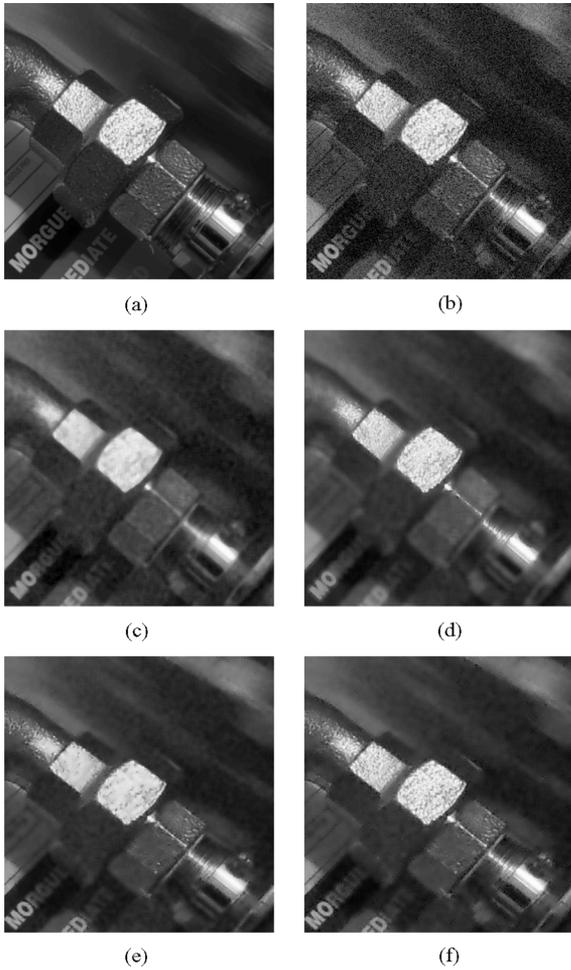


Fig. 13. (a) Original image and (b) the noisy version. (c) Median filter and (d) the adaptive Wiener filter. (e) Filtered image in IS and (f) the filtered image in LS.

TABLE II
MSE, MAE AND SNR MEASURES ASSOCIATED WITH “LENA”
IMAGE. NOISE PARAMETER $\gamma = 0.3$, $\sigma_u = 0.05$

Image	MSE	MAE	SNR (dB)
Noisy	322.704	13.644	17.39
Adaptive Wiener	115.747	6.895	21.84
Filtered in image space	63.574	5.509	24.44
Filtered in light space	61.037	5.366	24.62

image. This is because, in IS, the adaptive Wiener filter, addressing the last noise region, is not an optimal filter. A more complex filtering technique is needed to handle this type of noise. However, in LS, noise in middle and high intensity range follows the same behavior with positive γ . Therefore, ANS filter results are satisfactory for bright regions of the image (high intensities).

B. Real Images

To demonstrate light-space filtering on real images, a series of 201 images of a standard test pattern (see Fig. 14) were acquired from a Nikon D70 SLR digital camera. The camera was

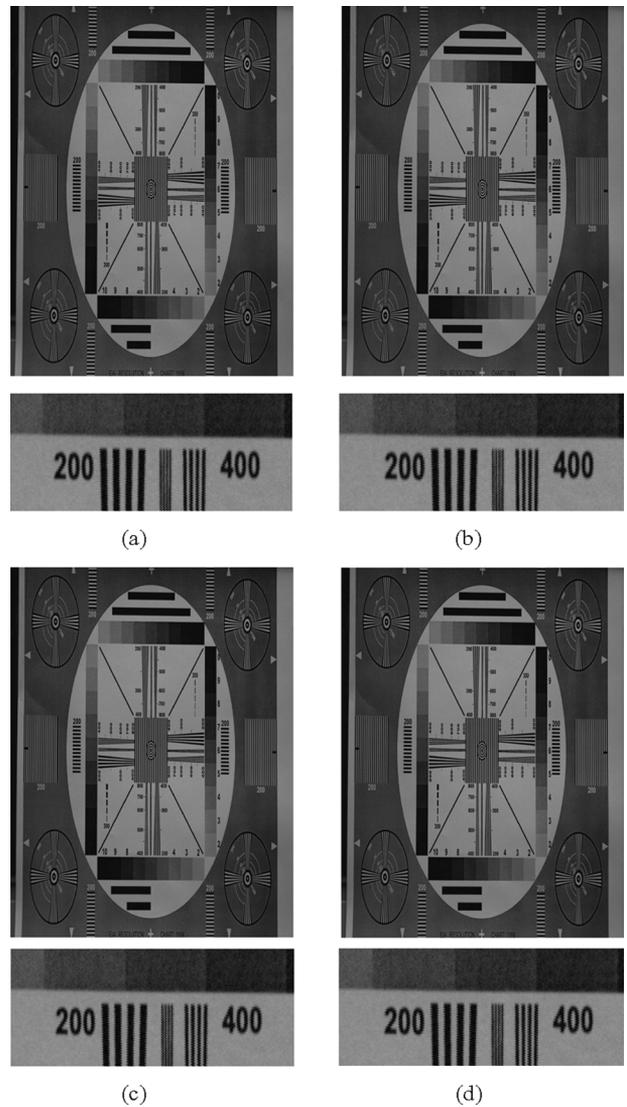


Fig. 14. This figure compares the (estimated) (a) noise-free, (b) noisy, (c) light-space filtered, and (d) image-space filtered images. The complete image is shown above, while a small region is enlarged below. Note that the filtered image is a substantial improvement over the noisy image, while retaining good edge definition. In uniform regions, the filtered image appears more uniform than the estimated noise-free image, although borders between regions with subtle intensity variations suffer somewhat.

set to ISO1600, F4.5 and a shutter speed of 1/80 second to generate sufficiently noisy images.² If one assumes that a zero-mean sensor noise model is accurate for this camera, then the mean of the images serves as an estimate of the noise-free image. It should be noted that with a finite sample size, and allowing for the possibility of camera shake and lighting variations, that this is still an approximation. The camera can be set to simultaneously output both 12-bit raw lightspace images and 8-bit image space images, making estimation of the CRF a simple matter. In order to apply our technique, the noise variance as a function of image intensity was plotted in light-space, and values for σ_u , γ , and σ_d were manually estimated (assuming a two-segment

²While these exposure settings may seem rather artificial, one can easily imagine settings such as these, for example, in the case of security cameras in low-light settings.

TABLE III
MSE RESULTS FROM D70 NOISY LIGHTSPACE IMAGES

	Noisy	Wiener	NANS	Final Combined
Min MSE	24.7193	11.4647	8.6748	8.9258
Average	33.9907	19.9292	17.0050	17.2690
Std Dev	12.7688	11.8942	11.7066	11.7235

MSE Ratios				
	Min	Wiener	NANS	Final Combined
Min	1.0000	0.4638	0.3509	0.3611
Average	1.0000	0.5618	0.4705	0.4788
Std Dev	0	0.0739	0.0901	0.0886

noise model) from the resulting graph, with values of 32.5, 0.2, and 4.4, respectively. Note that σ_u becomes 0.007 when normalized by 2^{12} , to allow direct comparison with results from the synthetic sequences. Similarly, in image space the parameters were found to be $\sigma_u = 44.6$, $\gamma = -0.47$, and $\sigma_d = 8$. Here, a two-segment curve has been used, but the nonconstant segment of the curve was found to have a negative slope. The deviation between the measured and theoretical values for γ can be attributed to internal signal processing within the camera.

Results from the sequence are shown in Table III and Fig. 14. While the improvement in MSE is not as dramatic as for the synthetic examples, note that the overall MSE for the noisy image is relatively small, and the NANS filter still outperforms the Wiener filter by a substantial margin. Since some systematic variation was found in the noisy images' MSE values (possibly indicating camera shake or lighting variation while the images were captured), a better figure of merit may be the ratio of post-filtered to prefiltered MSE values for the Wiener, NANS, and Final Combined filters over the image sequence: these results are also shown in Table III, and suggest superior performance from the Final Combined filter over the Wiener filter. The corresponding data for applying our method to image space images is shown in Table IV. While the performance of the two appear identical, this may be slightly deceptive as the test images did not span the full dynamic range of the camera, and the amount of data in the saturation region of the CRF is small, so we would expect similar performance. It is also worth noting that the estimated noise-free image still appears quite noisy, so the final combined performance, for both light space and image space, may still be better than the numbers would indicate. For a test sequence of 201 images, one would expect the effective σ_u in the mean image to be about 2.3. In tests on image regions that are expected to be uniform, the standard deviation in the mean images was found to be two- to three-times that of the combined final filtered images in general.

In summary, despite similar performance in both domains, LS restoration is preferable because 1) with a simpler noise model, LS denoising performance is as good or better than that in IS, and in bright regions it outperforms the second method, 2) LS denoising is computationally less demanding, due to fewer weight functions and filters applied, and 3) the

TABLE IV
MSE RESULTS FROM D70 NOISY IMAGESPACE IMAGES

	Noisy	Wiener	NANS	Final Combined
Min MSE	24.7193	11.2259	8.3871	8.3871
Average	33.9907	19.2732	16.5578	16.5578
Std Dev	12.7688	11.3060	11.4484	11.4484

MSE Ratios				
	Min	Wiener	NANS	Final Combined
Min	1.0000	0.4694	0.3507	0.3507
Average	1.0000	0.5624	0.4729	0.4729
Std Dev	0	0.0699	0.0914	0.0914

LS approach is robust to the variation in the location of the breakpoint partitioning the signal-independent region from the SDN region. In terms of computational timing, our filtering algorithm, implemented as nonoptimized MATLAB code running on a 512×512 image, takes about 12.7 s.

VI. CONCLUSION AND REMARKS

In this paper, we have described two software-based techniques to address CCD noise in digital images. In contrast to most restoration methods, we account for the nonlinear CRF in our restoration algorithm. These techniques, which are based on the PTC, involve applying multiple filters on a single noisy image, and a smooth combination of the results of the individual filters. Simulation results for both methods show promising performance; uniform regions are smoothed, while abrupt changes such as edges are preserved. A comparison of the performance of the two methods demonstrates efficient noise removal performance for both methods. However, light space denoising enables the design of a simpler and more straightforward restoration algorithm and, thus, is preferred.

ACKNOWLEDGMENT

The authors would like to thank S. Mann and K. Plataniotis for providing the experimental setup and for their helpful comments on the work. They would also like to thank C. Manders for his help in acquiring the Nikon D70 test images and computing the CRF.

REFERENCES

- [1] S. Mann, *Intelligent Image Processing*. New York: Wiley, 2002.
- [2] —, "Comparametric equations with practical applications in quantitative image processing," *IEEE Trans. Image Process.*, vol. 9, no. 8, pp. 1389–1406, Aug. 2000.
- [3] J. R. Janesick, *Scientific Charge-Coupled Devices*. Bellingham, WA: SPIE, 2001.
- [4] T. Mitsunaga and S. Nayar, "Radiometric self calibration," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 1999, pp. 374–380.
- [5] G. Pavlovic and A. Tekalp, "Restoration in the presence of multiplicative noise with application to scanned photographic images," in *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing*, 1990, vol. 4, pp. 1913–1916.
- [6] G. E. Healey and R. Kondepudy, "Radiometric CCD camera calibration and noise estimation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 16, no. 3, pp. 267–276, Mar., 1994.

- [7] Y. Tsin, V. Ramesh, and T. Kanade, "Statistical calibration of CCD imaging process," in *Proc. IEEE Int. Conf. Computer Vision*, 2001, vol. 1, pp. 480–487.
- [8] H. Faraji, "Light Space Adaptive Removal of CCD Noise in Digital Images," M.S. thesis, Univ. Toronto, Toronto, ON, Canada, 2004.
- [9] K. A. Jain, *Fundamentals of Digital Image Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1989.
- [10] I. Pitas and A. N. Venetsanopoulos, "Order statistics in digital image processing," *Proc. IEEE*, vol. 80, no. 12, pp. 1893–1921, Dec. 1992.
- [11] T. K. Darwin and A. Alexander, "Adaptive noise smoothing filter for images with signal-dependent noise," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-7, no. 2, pp. 165–177, Feb. 1985.
- [12] F. Argenti, G. Torricelli, and L. Alparone, "Signal-dependent noise removal in the undecimated wavelet domain," in *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing*, 2002, vol. 4, pp. 3293–3296.
- [13] B. Aiazzi, L. Alparone, and S. Baronti, "A robust method for parameter estimation of signal-dependent noise models in digital images," in *Proc. IEEE Int. Conf. Digital Signal Processing*, 1997, pp. 601–604.
- [14] H. Faraji and J. Maclean, "Adaptive suppression of CCD signal-dependent noise in light space," presented at the IEEE Int. Conf. Acoustics, Speech, and Signal Processing 2004.
- [15] S. Mann, C. Manders, and J. Fung, "The light space change constraint equation (LCCE) with practical application to estimation of the projectivity + gain transformation between multiple pictures of the same subject matter," in *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing*, Apr. 2003, vol. 3, pp. 481–484.



Hilda Faraji (M'04) received the B.A.Sc. degree in electrical engineering from the University of Tehran, Tehran, Iran, 1997, the M.A.Sc. degree from the Edward S. Rogers, Sr., Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON, Canada, in 2004, and the M.A.Sc. degree from the University of Saskatchewan, Regina, SK, Canada.

She is currently with Array Systems Computing, Inc.



W. James MacLean (SM'05) received the B.A.Sc., M.A.Sc., and Ph.D. degrees from the University of Toronto, Toronto, ON, Canada, in 1986, 1989, and 1996, respectively.

He is an Assistant Professor in the Edward S. Rogers, Sr., Department of Electrical and Computer Engineering, University of Toronto. His research interests include image motion analysis, the application of statistical methods to problems in computer vision and acoustic data, and high-performance implementations of computer vision algorithms using reconfigurable hardware. He recently organized the First International Workshop on Spatial Coherence for Visual Motion Analysis in Prague.

He is a member of the Professional Engineers Ontario.