

Impact of multi-scale retinex computation on performance of segmentation algorithms

Zia-ur Rahman[†], Daniel J. Jobson[‡], Glenn A. Woodell[‡], Glenn D. Hines[‡]

[‡]NASA Langley Research Center, Hampton, Virginia 23681.

[†]College of William & Mary, Department of Applied Science, Williamsburg, VA 23187.

ABSTRACT

Classical segmentation algorithms subdivide an image into its constituent components based upon some metric that defines commonality between pixels. Often, these metrics incorporate some measure of "activity" in the scene, e.g. the amount of detail that is in a region. The Multiscale Retinex with Color Restoration (MSRCR) is a general purpose, non-linear image enhancement algorithm that significantly affects the brightness, contrast and sharpness within an image. In this paper, we will analyze the impact the MSRCR has on segmentation results and performance.

1. INTRODUCTION

In previous papers¹⁻³ we have described the multi-scale retinex with color restoration (MSRCR). The MSRCR is a general-purpose image enhancement algorithm that provides simultaneous dynamic range compression and (approximate) color constancy. Here, by dynamic range compression we mean the ability to capture images that have both very bright and very dark regions at the same time, *without* either saturating the bright regions or clipping the dark regions. And, by color constancy we mean that the perceived color of an object remains essentially the same regardless of the "color" of the illuminant. Both these capabilities are found in the human visual system but not in most digital cameras, with the result that whereas most humans can see and recognize objects under all kinds of imaging conditions, the same is not true for digital imagers or for algorithms that rely on images obtained with these digital imagers for object recognition and classification. However, these images could be processed with the MSRCR before the application of the classification algorithm to alleviate the impact of these shortcomings of digital imagers. While the MSRCR cannot affect saturated and clipped regions in an image, it will reduce the variations that are introduced into the images due to changes in the irradiance across the image due to changes in illumination geometry, and reflectance changes. We have previously shown that the application of the MSRCR to multi-spectral data before classification can improve the performance of multi-spectral classification algorithms.^{4,5}

In this paper, we will attempt to quantify the impact of the MSRCR on a simple segmentation algorithm: the QUADTREE. We will use the idea of *visual measures* (VM)^{1,6} to segment an image into regions that have "sufficient" brightness and contrast and those that do not meet this criteria. In order to accurately measure the impact of different sources of classification errors, we will use the synthetic image shown in Figure 1. Specifically, we will examine 4 sources of error:

1. changes in illuminant color;
2. changes in dynamic range due to the presence of shadows;
3. changes in spatial resolution due to blurring;
4. changes in scene characteristics due to additive noise.

Each one of these sources of errors affects the image in a unique manner, and hence, can cause the segmentation and classification algorithms to fail. Since the MSRCR, or more specifically the MSR is color constant, the impact of changes in the color of the illuminant should be nullified. Similarly, the ability of the MSRCR to bring out details in areas where the radiance profile has been impacted by shadows should alleviate segmentation errors

Contact: ZR: zrahman@as.wm.edu; DJJ: daniel.j.jobson@nasa.gov; GAW: glenn.a.woodell@nasa.gov; GDH: glenn.d.hines@nasa.gov;

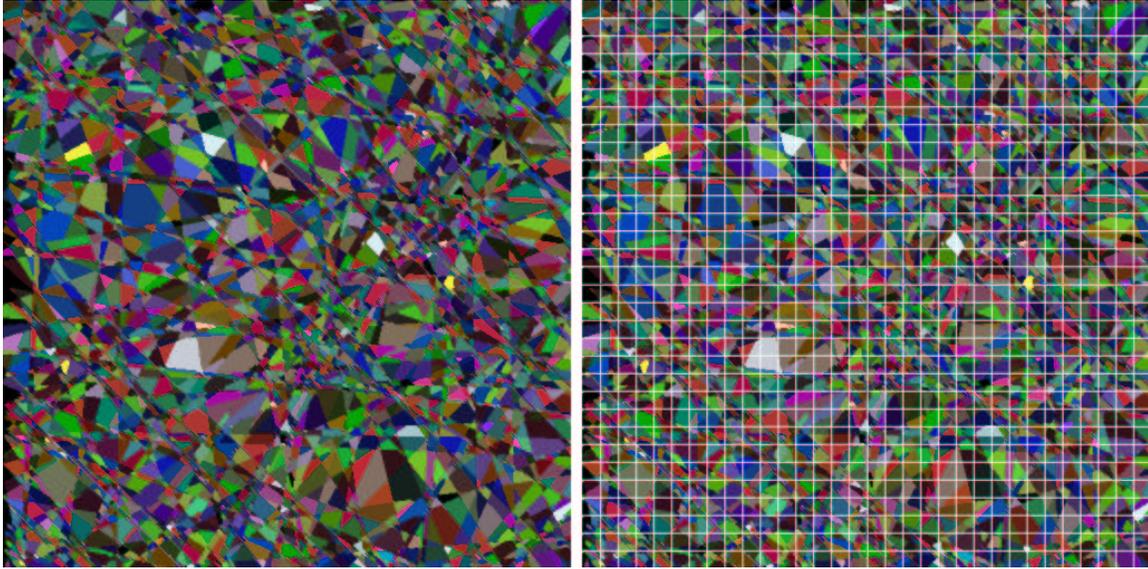


Figure 1. (a) The computer generated original image used to baseline the impact of the MSRCR on QUADTREE segmentation. All regions of this image meet the VM standards of sufficient brightness and contrast. The segmented image is shown in (b).

due to dynamic range variations. One of the unexpected properties of the MSRCR is an automatic sharpening of the original image. This can be related to the fact that the increased contrast that the MSRCR provides is translated into a perception of increased sharpness.⁷ The only source of error that the MSRCR may aggravate is that of additive sensor noise. In such a case, because the signal and noise are irrevocably combined in the output of the non-linear enhancement, the MSRCR may actually decrease the performance of the segmentation algorithm. Because noise adds random “details” to the original image, and the MSR enhances these details, this may be a situation where image classification may be less accurate when processed imagery is used in place of unprocessed imagery.

2. MULTI-SCALE RETINEX WITH COLOR RESTORATION (MSRCR)

The basic form of the multi-scale retinex (MSR) is given by

$$R_i(x_1, x_2) = \sum_{k=1}^K W_k (\log I_i(x_1, x_2) - \log [F_k(x_1, x_2) * I_i(x_1, x_2)]) \quad i = 1, \dots, \mathcal{S} \quad (1)$$

where index i references the i^{th} spectral band, (x_1, x_2) is the pixel location in Cartesian coordinates, and $*$ is the convolution operator. \mathcal{S} is the number of spectral bands— $\mathcal{S} = 1$ for grayscale images and $\mathcal{S} = 3, i \in R, G, B$ for typical color images. I is the input image and R is the output of the MSR process. F_k is the k^{th} (Gaussian) surround function, W_k is the weight associated with F_k , and K is the number of surround functions, or scales. The F_k are given as:

$$F_k(x_1, x_2) = \kappa \exp[-(x_1^2 + x_2^2)/\sigma_k^2],$$

where σ_k are the standard deviations of the Gaussian surrounds. The magnitude of σ_k controls the extent of the surround, smaller values of σ_k result in narrower surrounds. The MSR output is normalized by $\kappa = 1/(\sum_{x_1} \sum_{x_2} F(x_1, x_2))$. The MSR reduces to the single scale retinex (SSR) when $K = 1$, with the additional constraint that $W_1 = 1$. Figure 2 shows an input image, the output of the MSR, and the outputs when the different surround functions are applied to the original image. These are obtained by setting $k = 1$ and $W_k = 1.0$ in Equation 1. As is evident from Figure 2, SSR cannot attain visual realism. However, both the SSR and the MSR computations are completely color constant.²

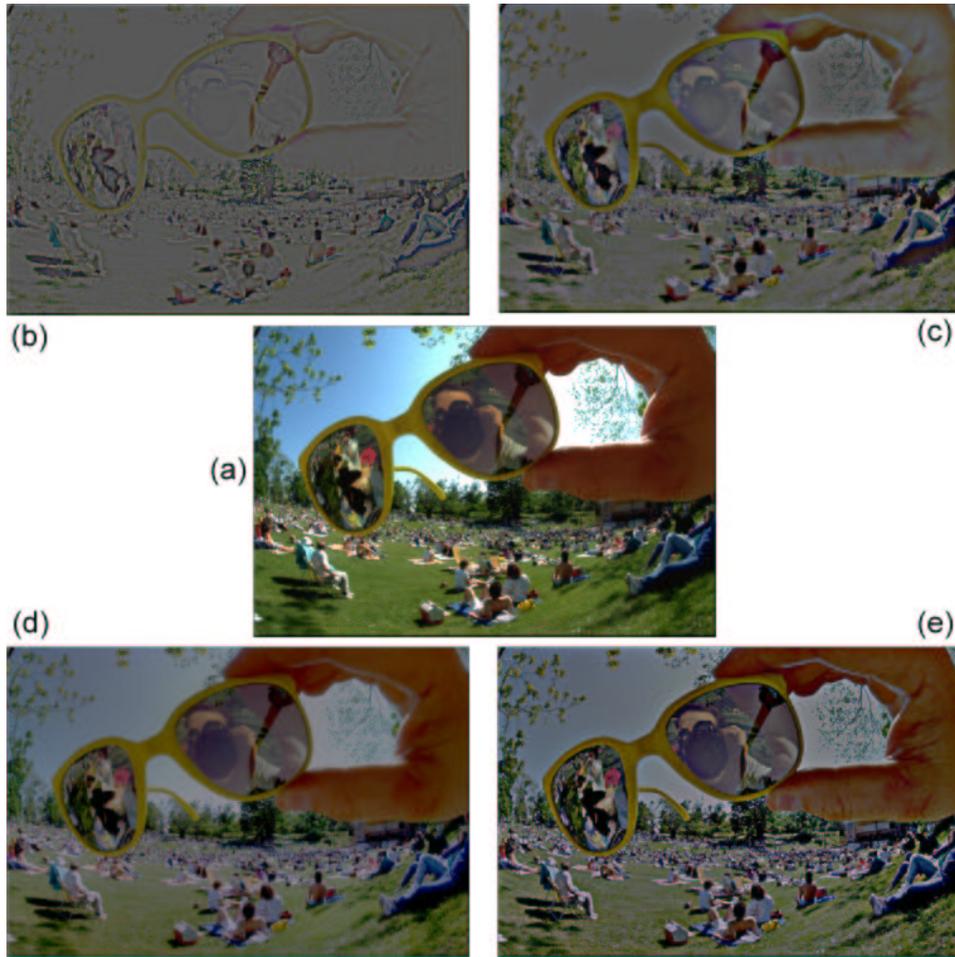


Figure 2. (a) The original input (b) Narrow surround $\sigma = 5$ (c) Medium surround $\sigma = 20$ (d) Wide surround $\sigma = 240$ (e) MSR output with $W_k = 1/3, k = 1, 2, 3$. The narrow-surround acts as a high-pass filter, capturing the fine detail in the image but at a severe loss of tonal information. The wide-surround captures the fine tonal information but at the cost of fine detail. The medium-surround captures both dynamic range and tonal information. The MSR is the average of the three renditions.

The general effect of MSR processing on images with regional or global gray-world violations is a “graying out” of the image either in specific regions or globally. This desaturation of color can, in some cases, be severe. We can, therefore, consider the computation that is needed to mitigate this desaturation as a color restoration (CR). The CR process should produce good color renditions for any degree of graying. In addition, the CR should preserve a reasonable degree of color constancy since that is one of the basic motivations for the MSR. However, color constancy is known to be imperfect in human visual perception, so some level of illuminant color dependency is acceptable, provided it is much lower than the physical spectrophotometric variations. Ultimately this is a matter of image quality, and color dependency is tolerable to the extent that the visual defect is not visually too strong.

Our starting point for the CR process is the foundations of colorimetry⁹ even though it is often considered to be in direct opposition to color constancy models and is felt to describe only the so-called “aperture mode” of color perception, i.e. restricted to the perception of color lights rather than color surfaces.¹⁰ The reason for this choice is simply that it serves as a foundation for creating a relative color space and in doing so uses ratios that are less dependent on illuminant spectral distributions than raw spectrophotometry. We compute a color

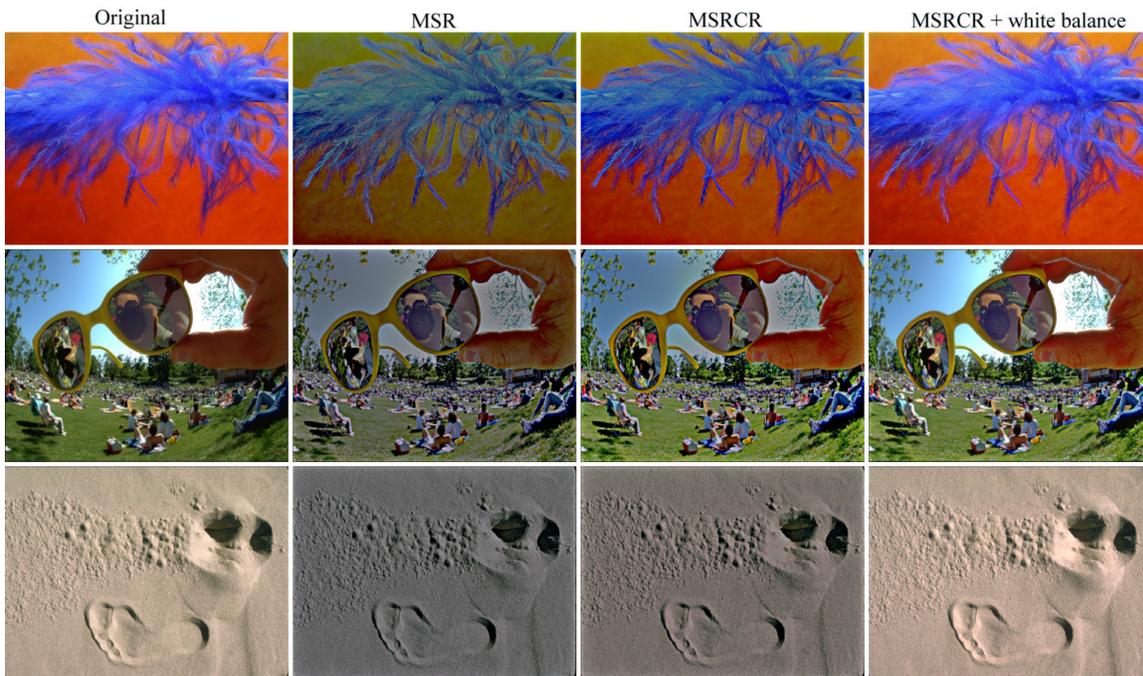


Figure 3. Scenes that violate the gray-world assumption, and the MSR, MSRCR and MSRCR with white-balance⁸ outputs. Note that while all the processed outputs are sharper than the originals, the MSR output is considerably more desaturated than the MSRCR output which still shows some color loss. The MSRCR with white-balance corrects the latter problem.

restoration factor, α based on the following transform:

$$\alpha_i(x_1, x_2) = f \left(I_i(x_1, x_2) / \sum_{n=1}^{\mathcal{S}} I_n(x_1, x_2) \right), \quad i = 1, \dots, \mathcal{S}, \quad (2)$$

where $\alpha_i(x_1, x_2)$ is the color restoration coefficient in the i^{th} spectral band, \mathcal{S} is the number of spectral bands, I_i is the i^{th} spectral band in the input image, and $f(\cdot)$ is some color space mapping function. Combining the color restoration term in Equation 2 with the MSR given in Equation 1, gives the Multiscale Retinex With Color Restoration (MSRCR):^{*†}

$$R_i(x_1, x_2) = \alpha_i(x_1, x_2) \sum_{k=1}^K W_k (\log I_i(x_1, x_2) - \log [F_k(x_1, x_2) * I_i(x_1, x_2)]), \quad i = 1, \dots, \mathcal{S}. \quad (3)$$

The results of applying this transformation to images with significant “monochrome” areas are shown in Figure 3. It is noticeable that the color restoration term does not completely restore the bright colors that are in the original image—see middle row in Figure 3. This effect can be ameliorated by using a white balance process that is the subject of a current patent application.⁸ In essence, the white balance process ensures that bright areas in the original image do not get desaturated to middle gray. Barnard and Funt¹¹ hypothesized that MSR processing suffers from “colour bleeding at certain colour edges due to the local contrast enhancement.” Though we see this effect in computer rendered images with very sharp edge transitions, we have not observed it to be a major source of concern in the many thousands of images we have processed. Barnard and Funt also point out in¹¹ that this “is normally not noticeable in images of typical natural scenes.”

^{*}See³ for more details on implementation.

[†]MSRCR has been implemented in a commercial software package, PhotoFlair, available from TruView Imaging Company.

While we have called this additional computation a color “restoration,” depending upon the form of $f()$, this can be considered as a spectral analog to the spatial retinex computation. If $f() = \log()$, then Equation 2 becomes

$$\alpha_i(x_1, x_2) = \log \left(\mathcal{S}I_i(x_1, x_2) / \sum_{n=1}^{\mathcal{S}} I_n(x_1, x_2) \right), \quad i = 1, \dots, \mathcal{S}$$

and the internal form of the retinex computation and the color restoration computation is essentially the same. This mathematical and philosophical symmetry is intriguing since it suggests some underlying unifying principle between the two computations: both computations are contextual, highly relative, and nonlinear. We speculate that the visual representation of wide ranging scenes is a compressed mesh of contextual relationships even at the stage of lightness and color representation. This sort of information representation would certainly be expected at more abstract levels of visual processing such as shape information composed of edges, links, and the like but is surprising for a representation so closely related to the raw image.

3. QUADTREE SEGMENTATION BASED ON VISUAL MEASURES (VM)

Though the QUADTREE algorithm is well-known in image segmentation literature,¹² we will describe it briefly here, especially with reference to the VM segmentation criteria. The QUADTREE is a recursive operation where each node of the tree, starting at the root, is subdivided into 4 subnodes—hence the name QUADTREE—depending upon some segmentation criteria. The basic algorithm is as follows:

1. Set the maximum depth of the QUADTREE to Λ , and the current level $l = 0$.
2. Set the root image $G_\lambda = G$, where G is the $N_1 \times N_2$ input image, and λ is the root index.
3. Set the number of rows of the root image to $\mathcal{R} = N_1/2^l$ and the number of columns to $\mathcal{C} = N_2/2^l$.
4. Divide the root $\mathcal{R} \times \mathcal{C}$ image, $G_\lambda[m, n]$, $m = 0, \dots, \mathcal{R} - 1, n = 0, \dots, \mathcal{C} - 1$ into 4 disjoint images, $G_{\lambda,i}[m, n]$, $i = 1, \dots, 4$ such that:

$$\begin{aligned} G_{\lambda,1}[m, n] &= G_\lambda[m, n] & 0 \leq m < \mathcal{R}/2; & & 0 \leq n < \mathcal{C}/2 \\ G_{\lambda,2}[m, n] &= G_\lambda[m, n] & \mathcal{R}/2 \leq m < \mathcal{R}; & & 0 \leq n < \mathcal{C}/2 \\ G_{\lambda,3}[m, n] &= G_\lambda[m, n] & 0 \leq m < \mathcal{R}/2; & & \mathcal{C}/2 \leq n < \mathcal{C} \\ G_{\lambda,4}[m, n] &= G_\lambda[m, n] & \mathcal{R}/2 \leq m < \mathcal{R}; & & \mathcal{C}/2 \leq n < \mathcal{C} \end{aligned}$$

5. Set $i = 1$.
6. For subimage $G_{\lambda,i}$, compute the local brightness—mean—and contrast—standard deviation.
7. If the computed brightness and contrast values exceed a threshold, then the region has sufficient brightness and should be segmented further. Set the root node to the current node and go to Step 3.
8. Otherwise the particular region has insufficient contrast and brightness. Go to the next region, and go to Step 6.
9. If $i = 4$, or $l = \Lambda$, i.e. all the subnodes have been checked, or the maximum depth has been reached, then go to the next higher level and go to the next node at that level. Go to Step 6.
10. If all nodes have been traversed, stop.

The computed brightness and contrast are adjusted at each level since the size of the node image is a function of the depth. The VM⁶ use a combination of the regional statistics in an image to define whether a region in the image has sufficient brightness and contrast. We have modified that approach to use as the criterion for segmentation. An example of using VM to perform QUADTREE segmentation is shown in Figure 1(right). The grid overlaying the image shows the final segmentation results: each block represents a region of the image where the contrast and brightness satisfied the sufficiency criterion.

4. IMPACT OF MSRCR PROCESSING ON SEGMENTATION

In this section, we will examine the impact of the MSR process on the performance of QUADTREE segmentation. When needed, we will augment the results from QUADTREE segmentation with other case appropriate segmentations to provide a fuller picture of the impact of the MSR on segmentation as a whole.

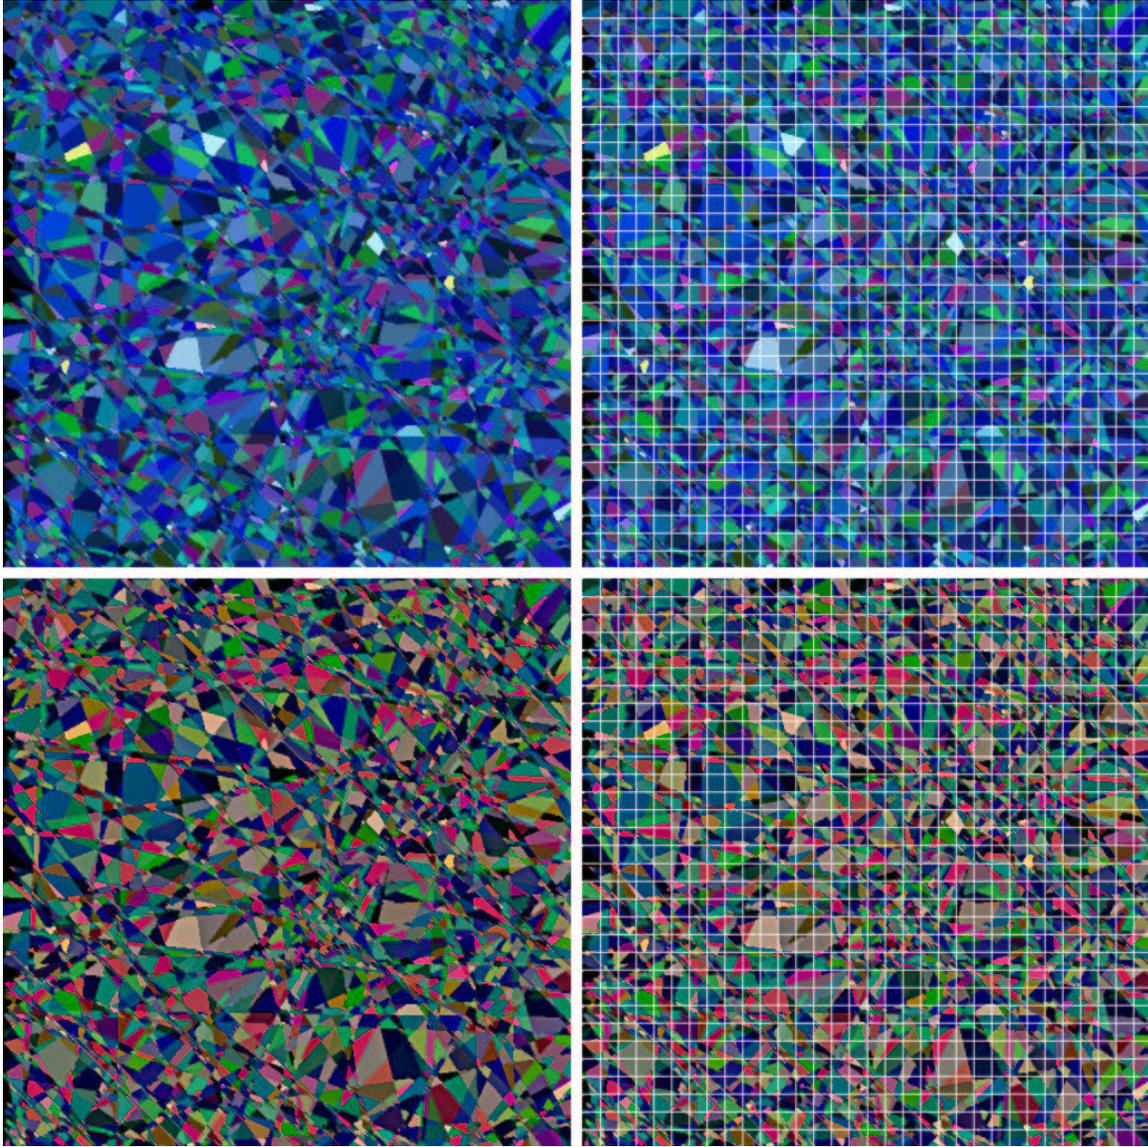


Figure 4. Top row: (left) image acquired under (simulated) blue illumination source; (right) QUADTREE segmentation. Bottom row: (left) MSR processed image; (right) QUADTREE segmentation on MSR processed images.

4.1. Illuminant color

We begin our analysis by examining the impact of changing the illuminant color on the QUADTREE segmentation. Figure 4 shows the image shown in Figure 1 when acquired under simulated blue light. While the overall color-tone of the image changes, making it *appear* bluer, the overall brightness and contrast within the image is affected only slightly. The QUADTREE segmentation classifies every region of the blue-shifted and MSR processed image to have sufficient brightness and contrast. In this particular case, a maximum depth $\Lambda = 8$ and smallest block-size of 50 were used for QUADTREE segmentation.

If, however, a different segmentation approach were under consideration, then the MSR processed image may have been more amenable to better segmentation. For instance, borrowing from the remote-sensing community, suppose that we were going to segment pixels based upon the relative strengths of response between different color bands: All those pixels where $I_r(x_1, x_2)/I_b(x_1, x_2) > 2.5$ are set to 255 and the rest to 0. Table 1 presents

Image	$I_r(x_1, x_2)/I_b(x_1, x_2) > 2.5$	$I_r(x_1, x_2)/I_b(x_1, x_2) < 2.5$	\mathcal{F}
Original	53202	208492	1.00
Blue-shifted	4191	257953	0.33
MSR processed	26400	235744	0.62

Table 1: Impact of MSR processing on classification based on multi-band ratio.

the result of performing this operation on the original image shown in Figure 1, the blue-shifted image shown top-left in Figure 4, and the MSR processed image shown bottom-left in Figure 4. The last column in Table 1, \mathcal{F} is a fidelity metric that measures similarity between the original image and the image under consideration.

$$\mathcal{F}(I_1, I_2) = 1 - \frac{\frac{1}{N_1 N_2} \sum_{x_1=0}^{N_1-1} \sum_{x_2=0}^{N_2-1} (I_1(x_1, x_2) - I_2(x_1, x_2))^2}{\frac{1}{N_1 N_2} \sum_{x_1=0}^{N_1-1} \sum_{x_2=0}^{N_2-1} I_1(x_1, x_2)^2}$$

Clearly the fidelity between an image and itself is 1. Thus, we are not only trying to measure the percentage of the number of classified pixels that also met classification criteria, but also the number of pixels that got classified similarly to the original image. As is clearly evident from Table 1, the MSR processing not only improved the total number of classified pixels when compared with the blue-shifted image, but the accuracy of the classification was also two-fold better.

4.2. Shadows

In many remote-sensing applications, the color of the illuminant is not an area of concern. However, changes in illumination due to the presence of clouds, and the relative movement of the Sun and Earth are of significant concern in the design and implementation of classification algorithms. We have shown previously,^{4,5} that MSR processing can lead to an improvement in classification accuracy. Here, we will examine this in the context of QUADTREE segmentation. Figure 5 shows the one-dimensional horizontal irradiance profile $\mathcal{M}(x_1)$ that we have used to model the shadow. The shadowed image L_s is generated by setting $L_s(x_1, x_2) = L(x_1, x_2)\mathcal{M}(x_1)\mathcal{M}(x_2)$. \mathcal{R} refers to the ratio between the maximum and minimum value in the profile, and hence determines the depth of the shadow. Normal shadows have $2.0 \leq \mathcal{R} \leq 4.0$.²

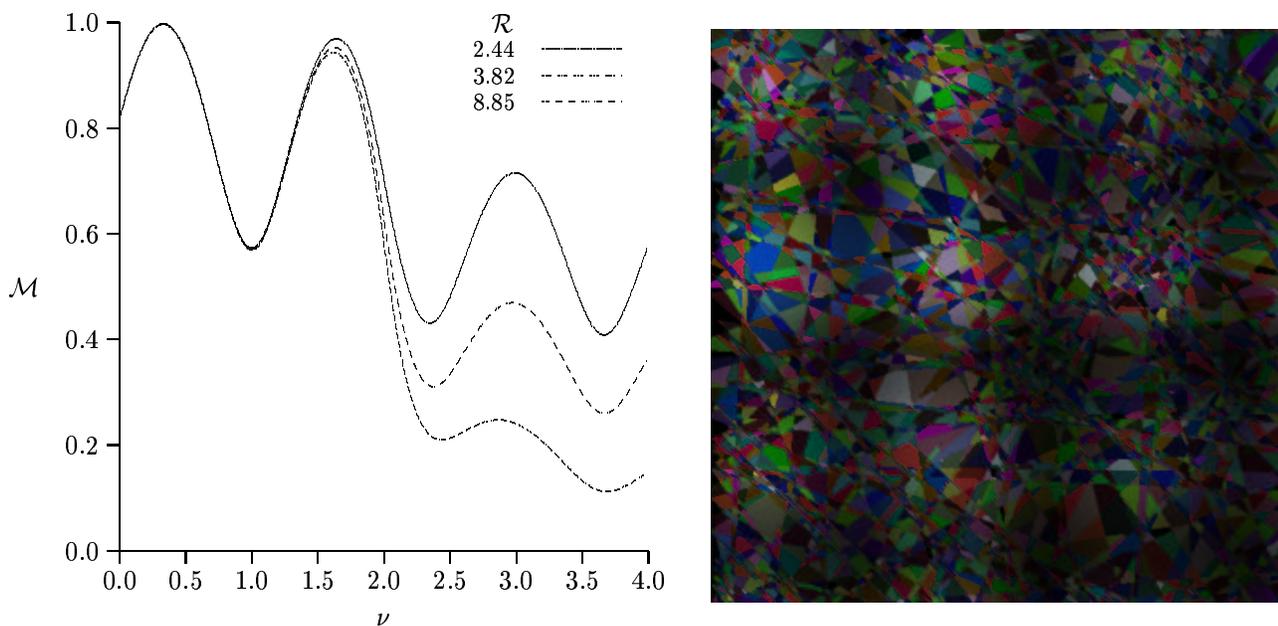


Figure 5: (a) Shadow profile (b) shadow profile with $\mathcal{R} = 2.44$ applied to the original image shown in Figure 1.

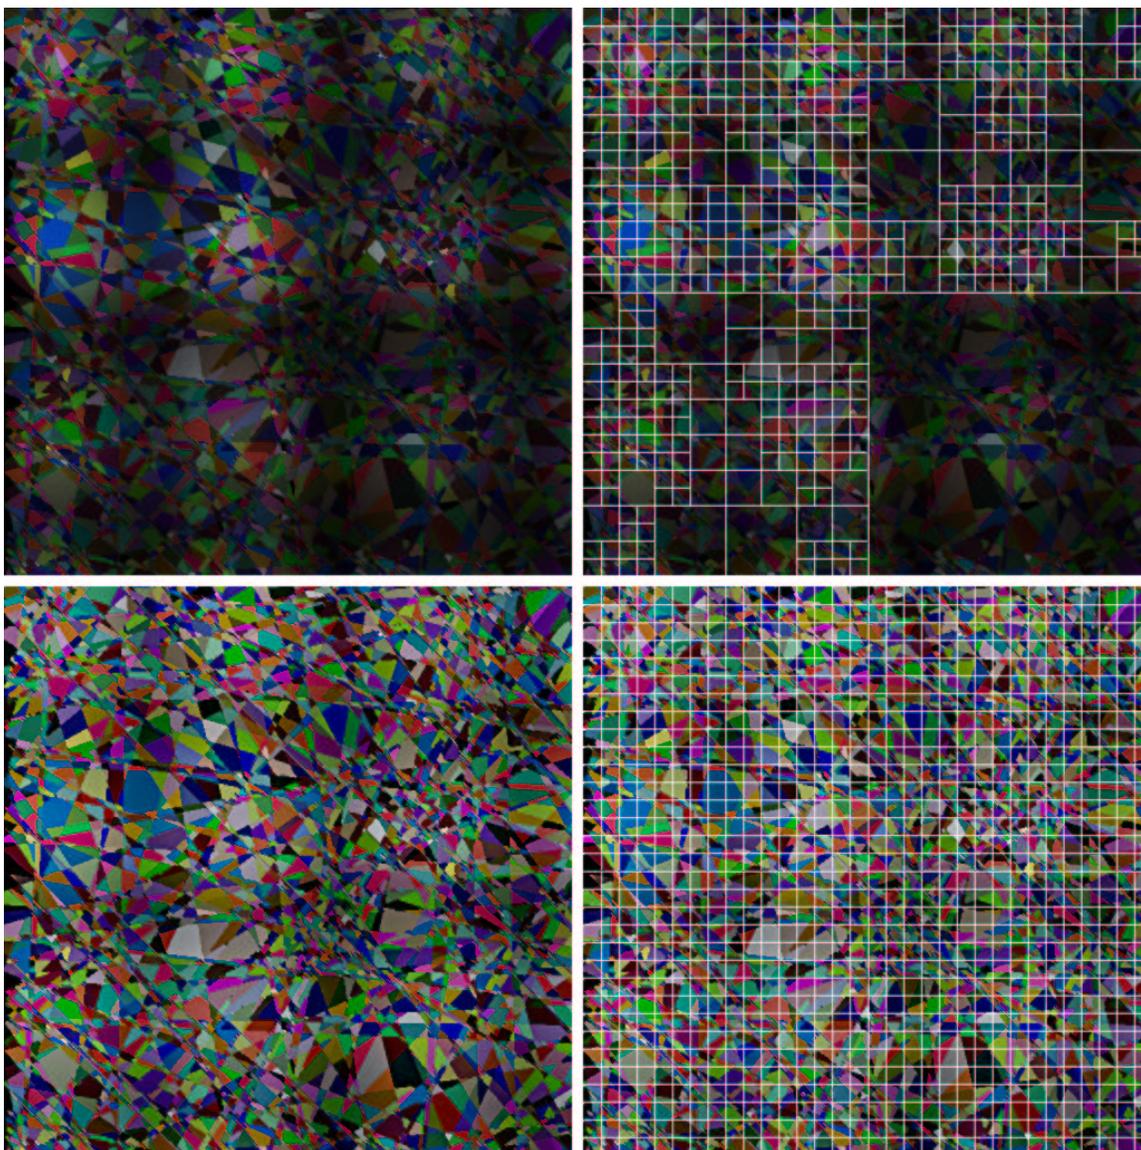


Figure 6. (top-row) Shadowed image and its QUADTREE segmentation (bottom-row) MSR processed shadowed image and its QUADTREE segmentation. The MSR processing (almost) completely removes the impact of shadows.

Figure 6 shows QUADTREE segmentation of the shadowed image (top-row) and the MSR processed shadowed image (bottom-row). For the shadowed image, QUADTREE segmentation determines that several areas do not satisfy the sufficient brightness and contrast criteria: Of the 1024 possible such regions, only 488 are classified as having sufficient brightness and contrast. The main reason for this lack is the bottom right quadrant of the image that fails the condition and is responsible for 256 out of 536 failures. The MSR processed image, however, completely restores the brightness and contrast profile, resulting in 100% of the regions being accurately classified! These results are summarized in Table 2.

4.3. Optical blurring

In order to measure the impact of optical blurring on QUADTREE segmentation, we blurred the image shown on the left in Figure 1 by convolving with a Gaussian blurring function that has the spatial-frequency domain

Image	\mathcal{P}_s	\mathcal{F}
Shadowed	47.66	0.69
MSR processed	100.00	0.95

Table 2. Impact of MSR processing on containing shadows. The fidelity \mathcal{F} is with respect to the original image shown in Figure 1. \mathcal{P}_s is the percentage of total regions that meet the sufficient brightness and contrast criteria.

representation given by

$$G(\nu, \omega) = \exp \left[-\frac{(\nu^2 + \omega^2)}{\rho^2} \right],$$

where (ν, ω) are spatial frequencies, and ρ is a parameter that controls the width of the Gaussian in the spatial-frequency domain. Heavier blurring is achieved with smaller values of ρ . This blurring function allows us to measure the impact of different levels of blurring on the segmentation process. In particular, we look at two cases: (1) $\rho = 0.05$ represents extremely heavy blurring and may correspond to the case where there is heavy fog or rain in the scene; (2) $\rho = 0.50$ represents light blurring and models the response of most image-capture devices. The $\rho = 0.50$ blurring also corresponds to the case where the only impact on the acquired image is of the lens itself and not of other atmospheric conditions. Figure 7 shows the blurred images, their QUADTREE segmentation and the corresponding MSR processed images and their QUADTREE segmentation. The numerical results are summarized in Table 3. As can be seen from the figure, QUADTREE segmentation based on brightness and contrast is fairly impervious to blurring: there is no impact on segmentation for light blurring and the image has to be blurred extremely heavily before any impact is evident at all. However, in either the light or the heavy blurring case, the MSR processed image has better contrast and sharpness than the blurred images. The fidelity information in Table 3 is quite interesting as well. In both cases, the fidelity of the blurred image is (very) slightly higher than the fidelity of the MSR processed image. This is easily explainable because the overall impact of the MSR process is to increase brightness, contrast, and sharpness of the image data. Hence, even when the

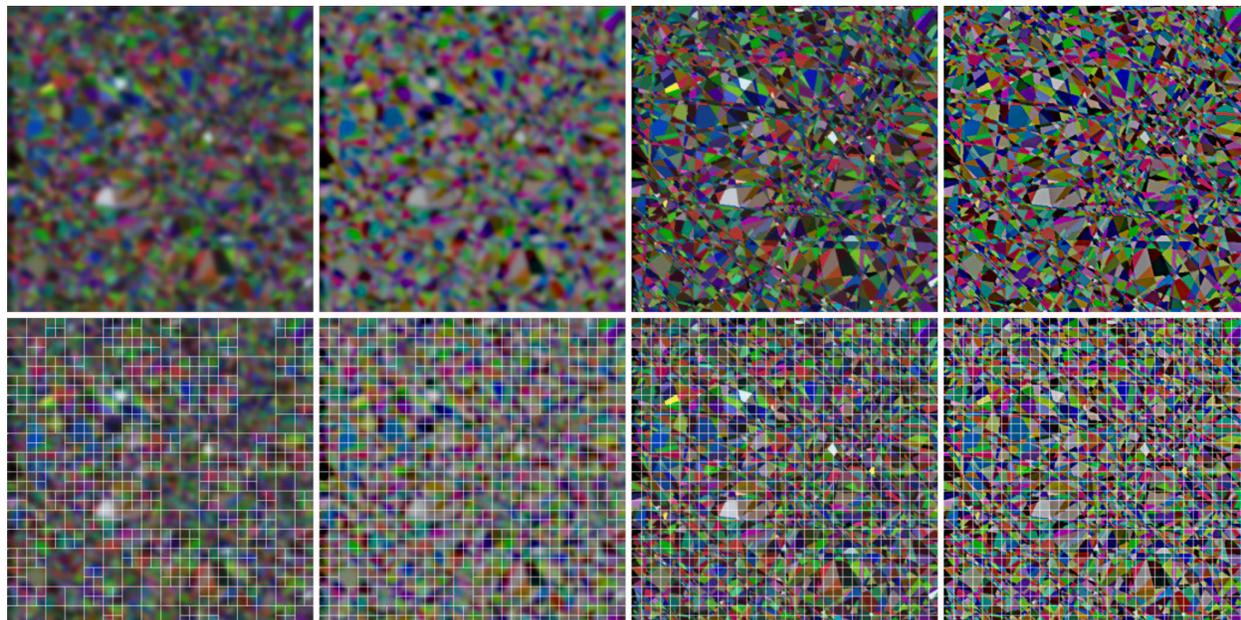


Figure 7. (first-column) Heavily blurred image with $\rho = 0.05$ and its QUADTREE segmentation; (second-column) MSR processed heavily blurred image and its QUADTREE segmentation; (third-column) slightly blurred image with $\rho = 0.50$ and its QUADTREE segmentation (last-column) MSR processed slightly blurred image and its QUADTREE segmentation. The MSR processing (almost) completely removes the blur for the lightly blurred case and improves the overall resolution for the heavily blurred case.

Image	\mathcal{P}_s	\mathcal{F}
Blurred with $\rho = 0.05$	56.64	0.83
MSR processed	92.19	0.82
Blurred with $\rho = 0.50$	100.00	0.99
MSR processed	100.00	0.95

Table 3. Impact of MSR processing on blurred images. The fidelity \mathcal{F} is with respect to the original image shown in Figure 1. \mathcal{P}_s is the percentage of total regions that meet the sufficient brightness and contrast criteria.

image has no shadows, blurring, or other artifacts, the MSR process will produce an image that is different in spectral distribution of the channels than the original image: a certain amount of change is the expected result of MSR processing.

4.4. Additive noise

The last error source that we address is that of additive noise, specifically, white Gaussian noise. The process is modeled as

$$I(x_1, x_2) = L(x_1, x_2) + N(x_1, x_2)$$

where L is the scene, N is the noise, and I is the image under consideration. We can control the amount of noise by defining the root-mean-square (RMS) signal-to-noise ratio (SNR) as $K\sigma_l^2/\sigma_n^2$, where σ_l and σ_n are the standard deviations of L and N respectively. Assuming Gaussian distributions for the scene and the noise, these σ represent the power in the signal and noise. The K in the equation is linear radiance-to-signal conversion constant and can be assumed here to be 1, without loss of generality.^{13, 14}

We will consider two cases of noise: $SNR = 1$ and $SNR = 10$. Both these SNRs represent noisy images since only noisy images are of interest to this analysis. The case where $SNR = 1$ represents a very heavy noise condition and the case where $SNR = 10$ represents a moderate to low SNR. Since the impact of additive noise often appears as increased detail in the image, we use a different computer generated scene for this simulation. This scene has large areas of constant intensity where the presence of noise can be easily seen and where the impact of noise is most visible. Figure 8 shows the impact of the noise on the MSR process as well as on the QUADTREE segmentation. Since the MSR increases local contrast, the overall result is to increase the appearance of noise strength. However, due to the non-linear nature of the noise, it is difficult, if not impossible, to mathematically characterize the impact of this process. Jobson, et al., in a companion paper, have attempted to analyze the impact of the enhancement on the noisy image.¹⁵

The overall segmentation and fidelity results are shown in Table 4.

Image	Original	SNR = 10	SNR = 1	MSR		
				Original	SNR = 10	SNR = 1
\mathcal{P}_s	83.20	86.72	100.00	86.72	97.66	100.00
\mathcal{F}	1.00	0.93	0.49	0.95	0.71	0.27

Table 4. Impact of MSR processing on noisy images. \mathcal{P}_s is the percentage of regions classified as bright, and \mathcal{F} is the fidelity between the original image and the image specified by the column heading. The MSR process inherently increases brightness, contrast and sharpness so the number of regions classified as having sufficient brightness and contrast increases.

5. CONCLUSIONS

We have examined the impact of MSR processing on QUADTREE segmentation and some other segmentation algorithms. While the MSR aids segmentation in cases of image blur, illumination changes such as shadows, and changes in color of the illuminant, it hinders segmentation in the case of noisy images. In this paper we have concentrated on a relatively simple segmentation: that of defining regions as having sufficient brightness and contrast. A more telling evaluation would be to measure the impact of MSR processing on other classification

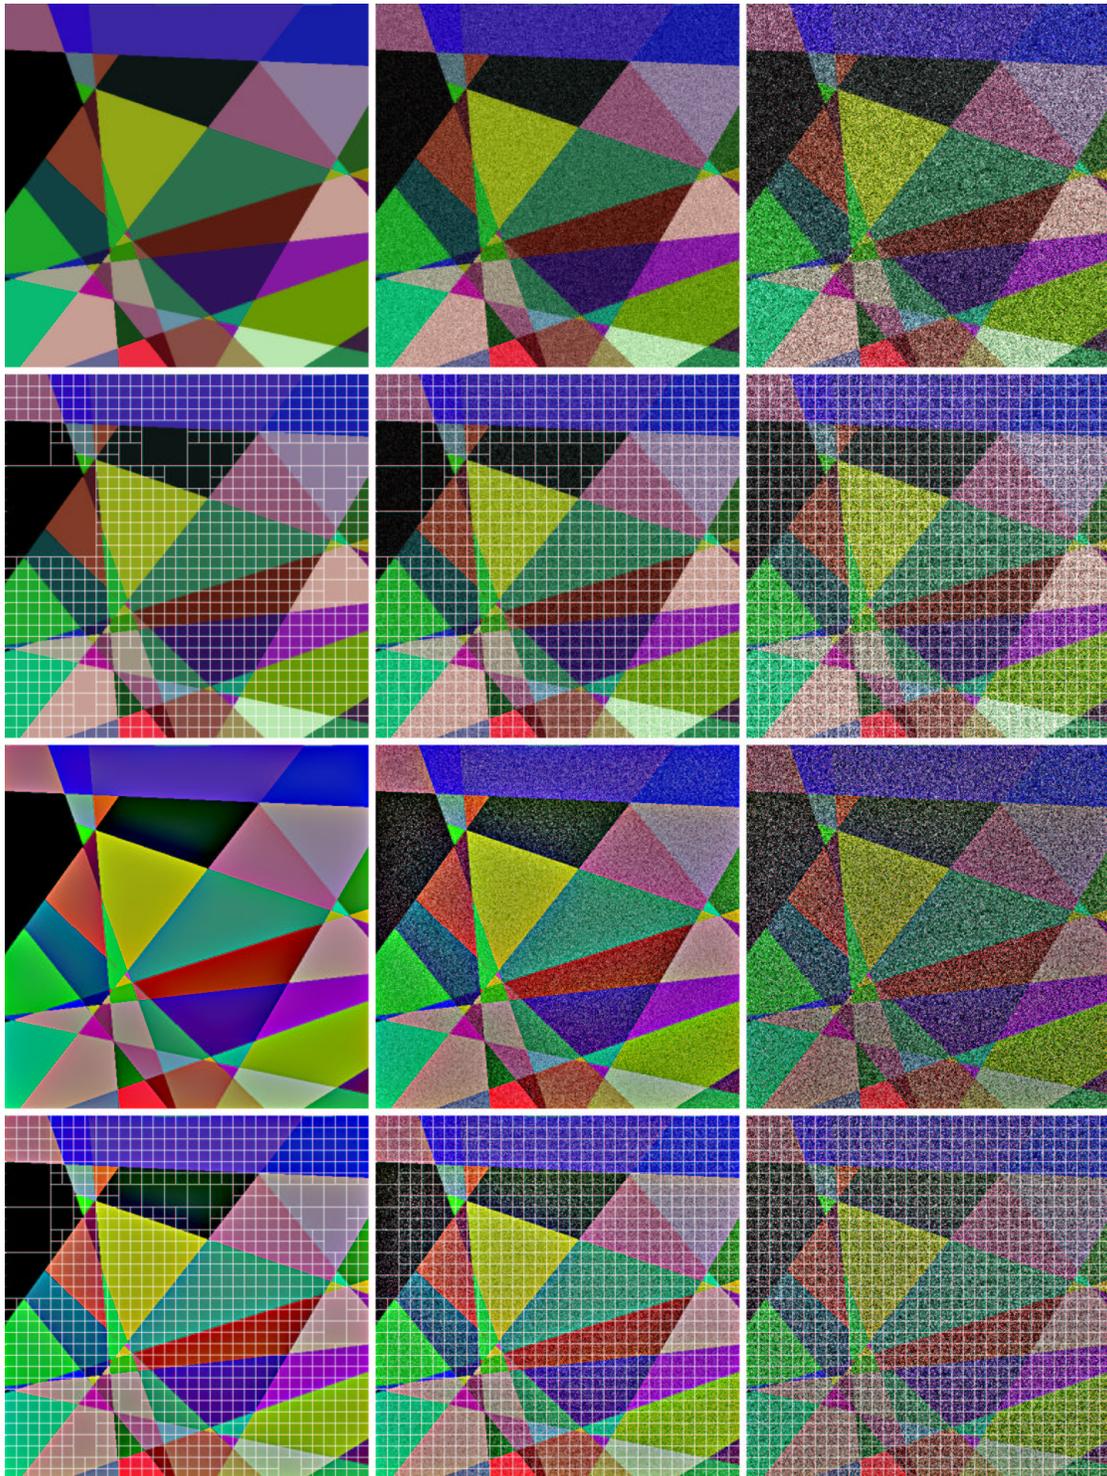


Figure 8. Figure 8 (top-row) shows the original noise-free image and the two images where noise has been added to the image. The second row shows the segmentation obtained when the contrast and brightness criteria QUADTREE is applied to these images. The third row shows the MSR processed versions of the three images, and the bottom row shows the QUADTREE segmentation for each case. The MSR process increases the overall noise and thus affects the segmentation negatively.

algorithms such as the well known K-means segmentation. Also, criterion other than brightness and contrast would be of interest. What the initial analysis shows is that there is considerable promise in using the MSR before segmentation. However, application of this algorithm to real-world scenes is a topic for future research.

Acknowledgments

The authors wish to thank the Synthetic Vision Sensors element of the NASA Aviation Safety Program for the funding which made this work possible. In particular, Dr. Rahman's work was supported under NASA cooperative agreements NCC-1-01030 and NNL04AA02A.

REFERENCES

1. Z. Rahman, D. J. Jobson, and G. A. Woodell, "Retinex processing for automatic image enhancement," *Journal of Electronic Imaging* **13**(1), pp. 100–110, 2004.
2. D. J. Jobson, Z. Rahman, and G. A. Woodell, "Properties and performance of a center/surround retinex," *IEEE Trans. on Image Processing* **6**, pp. 451–462, March 1997.
3. D. J. Jobson, Z. Rahman, and G. A. Woodell, "A multi-scale Retinex for bridging the gap between color images and the human observation of scenes," *IEEE Transactions on Image Processing: Special Issue on Color Processing* **6**, pp. 965–976, July 1997.
4. B. D. Thompson, Z. Rahman, and S. K. Park, "A multi-scale retinex for improved performance in multi-spectral image classification," in *Visual Information Processing IX*, Proc. SPIE 4041, 2000.
5. B. D. Thompson, Z. Rahman, and S. K. Park, "Retinex preprocessing for improved multi-spectral image classification," in *Visual Information Processing VIII*, Proc. SPIE 3716, 1999.
6. D. J. Jobson, Z. Rahman, and G. A. Woodell, "The statistics of visual representation," in *Visual Information Processing XI*, Z. Rahman, R. A. Schowengerdt, and S. E. Reichenbach, eds., pp. 25–35, Proc. SPIE 4736, 2002. Invited paper.
7. W. F. Schreiber, "Psychophysics and the improvement of television image quality," *SMPTE Journal* **93**(8), pp. 717–725, 1984.
8. G. A. Woodell, D. J. Jobson, and Z. Rahman, "Method of improving a digital image having white zones," *U.S. Patent Application 20030026494*, 2001.
9. W. D. Wright, *The Measurement of Colour*, Hilger and Watts, London, second ed., 1958.
10. P. Lennie and M. D. D'Zmura, "Mechanisms of color vision," in *CRC Critical Reviews of Neurobiology*, vol. 3, pp. 333–400, 1988.
11. K. Barnard and B. Funt, "Investigations into multi-scale retinex," in *Colour Imaging: Vision and Technology*, pp. 9–17, John Wiley and Sons, New York, NY, 1999.
12. A. Rosenfeld and A. C. Kak, *Digital Picture Processing*, vol. 2, Academic Press, Orlando, FL, second ed., 1982.
13. F. O. Huck, C. L. Fales, and Z. Rahman, "Information theory of visual communication," *Philosophical Transactions of the Royal Society of London A* **354**, pp. 2193–2248, Oct. 1996.
14. F. O. Huck, C. L. Fales, and Z. Rahman, *Visual Communication: An Information Theory Approach*, Kluwer Academic Publishers, Norwell, MA, 1997.
15. D. J. Jobson, Z. Rahman, G. A. Woodell, and G. D. Hines, "The automatic assessment and reduction of noise using edge pattern analysis in nonlinear image enhancement," in *Visual Information Processing XIII*, Z. Rahman, R. A. Schowengerdt, and S. E. Reichenbach, eds., Proc. SPIE 5438, 2004.