

Image Processing for Human Understanding in Low-visibility

ABSTRACT

Low-visibility conditions for navigation of vehicles are a frequent occurrence. Driving at night, in blizzards, in sand storms, or in fog form an obvious set of challenging conditions. Remote operation of unmanned vehicles through a camera image provides a similar difficulty. Gathering intelligence from satellite imagery can similarly benefit from improved visibility. Advanced image processing techniques (e.g., contrast enhancement or tone mapping) purport to improve the perceptual quality of images that lack the contrast or color depth perceived by the human visual system (HVS). Applying such an algorithm intelligently to these low-visibility conditions gives us the ability to provide a perceptually usable assisted-vision system.

One premiere method for perceptual enhancement emerged from Retinex theory (Land, 1977; McCann, 2004). The key observation was that perceived color and intensity of a region in an image depend on not only inherent color and intensity, but also color and intensity of surrounding regions and on lighting. This property of the HVS serves as the basis for Retinex theory and numerous applications of it that process digital images to adjust color and intensity for the perceptual advantage of a human observer.

Original Retinex-based algorithms assumed that an image was underexposed; extensions have enabled contrast enhancement in overexposed images to be darkened for human perception. We contribute flexibility to Retinex processing by automatically determining multiple intensity levels from which brightening and darkening of imagery may be performed. We identify spatial areas in each color channel with similar intensity and therefore low local contrast. Pixels close to these intensity levels are treated with extra emphasis to reveal detail that may have previously been hidden to the eye. Each identified level is weighted against the others based on the prevalence. We show that our adaptation improves local contrast in a varied set of test cases which would benefit from perceptual enhancement.



Figure 1: A low-visibility image and our enhanced version

INTRODUCTION

Situations in which human performance is limited by low visibility are easy to conceive. Figure 1 shows an obvious example of driving on a rural road during a snow storm. Improved visibility of satellite imagery and live video from unmanned vehicles (notably undersea vehicles) further illustrate challenges of low visibility in video imagery. An ideal algorithm to process images requires no information about the sensor from which the image was acquired. This enables the algorithm to work for the widest variety of images and sensors. The result has obvious implications for both consumer vehicles and military situation awareness, as it greatly increases the distance the drive could effectively see. This ideal algorithm would further work automatically on a wide variety of low-

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE

OCT 2011

2. REPORT TYPE

3. DATES COVERED

00-00-2011 to 00-00-2011

4. TITLE AND SUBTITLE

Image Processing for Human Understanding in Low-visibility

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

**Naval Research Laboratory, Information Technology Division, 4555
Overlook Avenue SW, Washington, DC, 20375**

8. PERFORMING ORGANIZATION
REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT
NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

ASNE Human Systems Integration Symposium, 25-27 Oct 2011, Vienna, VA.

14. ABSTRACT

Low-visibility conditions for navigation of vehicles are a frequent occurrence. Driving at night, in blizzards, in sand storms, or in fog form an obvious set of challenging conditions. Remote operation of unmanned vehicles through a camera image provides a similar difficulty. Gathering intelligence from satellite imagery can similarly benefit from improved visibility. Advanced image processing techniques (e.g., contrast enhancement or tone mapping) purport to improve the perceptual quality of images that lack the contrast or color depth perceived by the human visual system (HVS). Applying such an algorithm intelligently to these low-visibility conditions gives us the ability to provide a perceptually usable assisted-vision system. One premiere method for perceptual enhancement emerged from Retinex theory (Land, 1977; McCann 2004). The key observation was that perceived color and intensity of a region in an image depend on not only inherent color and intensity, but also color and intensity of surrounding regions and on lighting. This property of the HVS serves as the basis for Retinex theory and numerous applications of it that process digital images to adjust color and intensity for the perceptual advantage of a human observer. Original Retinex-based algorithms assumed that an image was underexposed; extensions have enabled contrast enhancement in overexposed images to be darkened for human perception. We contribute flexibility to Retinex processing by automatically determining multiple intensity levels from which brightening and darkening of imagery may be performed. We identify spatial areas in each color channel with similar intensity and therefore low local contrast. Pixels close to these intensity levels are treated with extra emphasis to reveal detail that may have previously been hidden to the eye. Each identified level is weighted against the others based on the prevalence. We show that our adaptation improves local contrast in a varied set of test cases which would benefit from perceptual enhancement.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 9	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

visibility images without the need for “tuning” parameters (Funt et al., 2002). In this paper, we shall describe our efforts to build such an algorithm using Retinex theory as a basis. We show examples of successful application of our algorithm, cases that our algorithm fails to satisfactorily improve an image, and numerical performance of our algorithm.

Retinex theory is the name given to the complex processing done in the *retina* and *cortex* within the human visual system. According to this theory (Land, 1977; McCann, 2004), humans perceive color through a process of spatial comparisons. The visual system compares lightness within the three bands corresponding to short-, medium-, and long-wave photoreceptors (commonly referred to respectively as blue, green, and red cones). The spatial comparison of these three lightness values determines the color of a given region.

These observations provide a basis for enhancing the perceptual quality of images. A classic example from Land and McCann (1971) illustrates the difficulty of the problem that the human visual system is asked to solve. Take the example of a black cat in sunlight and a white cat in the shade; they could send the same amount of (achromatic) light to the eye. The surrounding context helps the visual system separate the lighting and the native color of the objects. This observation is the basis for optical illusions, such as the checker shadow illusion¹, in which the context fools the eye into incorrectly perceiving identical colors as different.

These same observations have practical use as the basis for algorithms to improve visibility in adverse conditions (Rahman et al., 2004), such as fog and other forms of atmospheric haze, night-vision imagery, and underwater imagery, as well as particle interference such as blizzards (as seen in Figure 1), sand storms, or rain. In comparison to frequency-domain filters, Retinex benefits from operating on pixel values rather than global basis functions. Compared to tools such as histogram equalization, Retinex-based processing can more strongly enhance local contrast, since it operates explicitly on the ratio of logarithms of those pixel values. These operators were found by Land, McCann, and their colleagues to accurately predict the visibility to the human eye.

We describe some strategies and advantages for Retinex implementations, present our formulation to overcome some of these deficiencies, apply an objective measure of contrast to analyze the results,

and show resulting images. We conclude with a discussion of future capabilities that would increase the value of our image enhancement algorithm.

RETINEX FORMULATIONS

Reading Land’s description of Retinex (Land, 1977), one would understand Retinex as a path-based algorithm. Land found that a measureable, physical correlate for lightness was scaled integrated reflectance. This integration (of log-values) occurs over multiple paths by which one might approach a region. These integrated values are averaged over each path and for each color receptor (cone). Thus the relationship between a region i and the surrounding areas j can be summarized by the equations and schematics in Figure 2.

$$R^A(i, j) = \sum_k \delta \log \frac{I_{k+1}}{I_k} \quad \bar{R}^A(i) = \frac{\sum_{j=1}^N R^A(i, j)}{N}$$

$$\delta \log \frac{I_{k+1}}{I_k} = \begin{cases} \log \frac{I_{k+1}}{I_k}, & \left| \log \frac{I_{k+1}}{I_k} \right| > \text{threshold} \\ 0, & \left| \log \frac{I_{k+1}}{I_k} \right| < \text{threshold} \end{cases}$$

Figure 2: Retinex path formulation (Land, 1986)

Ng and Wang (2011) insightfully summarize attempts to improve on this model. We build on their review to describe the background from which we began our research.

A variety of paths have been applied to approach a region, including linear, spiral, and Brownian motions. Such path-based formulations introduce parameters that need to be “tuned” in order to optimize the performance of the algorithm; this results in high computational complexity. Recursive methods replace the paths with a recursive matrix calculation. This yields improved efficiency. However, the unknown number of levels required in the recursion significantly influences the quality of the results. This again limits the computational efficiency of the approach.

The surrounding intensity (per channel or in luminance) of a region may also be captured through functional forms. A Gaussian kernel is a frequently-used operator in image processing; it performs better than other functions for this purpose as well. A kernel that uses not intensity but contrast comes closest to the results of our work, in that it can enhance from both lesser and greater intensity and thus enhance both under- and over-exposed imagery.

A variational formulation of Retinex models the perceptual principles with an energy function,

¹ http://web.mit.edu/persci/people/adelson/checkershadow_illusion.html

solving for the optimal solution for some combination of reflectance and illumination. A variational model can allow for more flexible local control but also loses computational efficiency from the optimization. In addition, the customary issues with non-linear optimization (instability growing with the number of parameters, unclear rate of convergence, local minima) may limit the performance of such algorithms. Similar approaches include solving partial differential equations or Poisson equations. The latter assumes that illumination changes smoothly, while reflectance remains piece-wise constant. These assumptions are required for regularization (reducing over-fitting by imposing further constraints on the solution).

Our primary goal was to increase the flexibility of Retinex-based processing by designing an algorithm that could determine which regions or pixels to lighten and which to darken in an image, then simultaneously process these pixels. Our approach to the algorithm builds on the use of functional forms; we maintain the use of convolution with a Gaussian function to achieve this multi-centered extension to Retinex-based processing.

ADAPTED RETINEX ALGORITHM

Our processing pipeline consists of a pre-processing phase, conduct of our Retinex-based enhancement algorithm, and a post-processing phase. In this section, we describe each of these phases.

Pre-processing

Our innovation in the pre-processing stage consists of identifying “key” image intensities that will drive the adaptation of the remaining image intensities. We define these “centering” points for the Retinex-based algorithm as the intensities that have the greatest number of contiguous occurrences in the image. An image with both a dark shadow and a cloud would have both dark and light centers from these phenomena (respectively).

In order to calculate these centers, an application specific method was created which processes each color channel individually. It starts at every unmarked pixel and conducts a radial, breadth-first search in every direction. Each visited pixel is marked as having a comparable intensity if the pixel is unmarked and the difference in its intensity and the mean luminance of previously visited pixels was less than the standard deviation of intensities in that color channel divided by four. The result is many groups of contiguous pixels that have similar

intensities. The algorithm merges groups that may not be tangential but have a similar intensity. Next, the method eliminates groups from the list of centers if their number of pixels is less than 10% of the image. This insures that only important groups are selected as centers. The influence R of each center is created. This coefficient is determined by the number of pixels each center marked where a larger number of pixels corresponds to a larger influence, and the sum of all the influence coefficients is 0.75. Finally, the method adds the mean intensity of that channel to the center list to negate any effects of a poorly chosen center or lack of a center. The influence of this center is automatically 0.25. The processing described in the next section will run with each center, scale the output relative to that center’s importance, and sum the scaled outputs to compute the final pixel value.

We note one other pre-processing step that we take for purely practical implementation reasons: we pad the image to avoid interference from the cyclic properties of the Fourier transform, which we use to implement the convolution operator.

Adapted Retinex Processing

When expressed in its original form, Retinex-based processing exclusively would brighten an under-exposed image, but kernels that focus on contrast rather than lightness can enhance both over- and under-exposed images (Bertalmio et al., 2009). We introduce the same flexibility in our algorithm by selecting the “center” point from which the lightness values would be pushed away. Mathematically, this means our first step was to convert the multi-scale Retinex formulation (Rahman et al., 2004) from its original basis in Equation 1 to a basis that would enable this centering mechanism.

$$R_i(x, y) = \sum_{k=1}^K W_k \left[\ln \left(\frac{I_i(x, y)}{F_k(x, y) * I_i(x, y)} \right) \right]$$

Equation 1: Multi-scale Retinex formulation

Equation 1 operates on each color channel i (taking the values of red, green, and blue) and at three kernel sizes k (known as in Multi-scale Retinex as small, medium, and large). The familiar logarithm of ratios of pixel intensity to the average local intensity (embodied in the convolution operator, $*$) yields the Retinex-based adjustment of image intensities. The typical enhancement created by this formulation is shown in Figure 3. Input pixel intensity (horizontal axis) is mapped by the logarithm function (vertical), producing the image brightening described above.




Figure 3: Shape of enhancement function of Multi-scale Retinex algorithm

First, we express the Multi-scale Retinex equation in a new form, shown in Equation 2.

$$R_i(x, y) = \sum_{k=1}^K W_k \left[-\ln \left(1 + \frac{F_k(x, y) * I_i(x, y) - I_i(x, y)}{I_i(x, y)} \right) \right]$$

Equation 2: Re-expression of Multi-scale Retinex

When expressed in this form, the equation can be adapted in several key ways. First, it accounts for the difference between the pixel and its background, expressed as the numerator in the fraction. When the Multi-scale Retinex algorithm was re-expressed, it was assumed that the difference between the background and the pixel was positive. This difference must always be positive to always be in the domain of the natural logarithm function; thus, an absolute value is added. However, the sign of this difference determines the reflection of the function along the horizontal axis in Figure 3. The sign of the difference is therefore extracted out of the natural logarithm and made the coefficient of the natural logarithm term. Note that this slightly alters the numerical behavior of the algorithm from its original form; however, this alteration makes the equation more stable and fair. Now that the domain issues are resolved, the denominator of the fraction can be made the absolute value of the difference between the pixel and the value of a vertical asymptote to center the function on that vertical asymptote. A b^{th} root operation limits and adjusts the asymptotic behavior; a constant D ensures that all pixels are in the domain. Finally, we sum the output for each center, scaled by the importance coefficient R for that center. The resulting formulation appears in Equation 3; a potential two-center enhancement function appears in Figure 4.

$$AR_i(x, y) = \sum_{k=1}^K \sum_{n=1}^N W_k R_n \text{sgn}(I_i(x, y) - F_k(x, y) * I_i(x, y)) \ln \left(1 + \frac{|F_k(x, y) * I_i(x, y) - I_i(x, y)|}{\sqrt[b]{D + |I_i(x, y) - C_n|}} \right)$$

Equation 3: Formulation of the adapted Retinex-based processing algorithm




Figure 4: A potential two-center enhancement function using our adapted Retinex algorithm and compared to Figure 3. Here, two vertical asymptotes (“centers”) have been selected, and neither is at the darkest or lightest input intensity. Note that the difference in output range is due merely to parameters used to draw these graphs.

Post-processing

The output of the natural logarithm function must be scaled back to the digital image range of $[0, 255]$. We scale to a wider range of $[-40, 295]$ and then clamp the values to $[0, 255]$ in order to reduce the effect of outlier pixels (e.g. typical CCD noise) and focus the output on the newly contrasting image intensities. This does not sufficiently protect the algorithm from enhancing noise and compression artifacts. Thus we introduce a smoothing algorithm that works in concert with the centering described above.

Any pixels that were not selected as a center in any color channel have their values fixed during this smoothing step. This prevents useful information in the original image from being blurred. Next, a similar selection of “centers” is calculated by considering the distance in each color channel *and* the Euclidean distance in color space. This criterion targets pixels that had an intensity change but no significant color change; we found such pixels often corresponded to compression artifacts and noise. Pixels not selected have their values fixed with no smoothing applied. The remaining pixels are converted to the YIQ color space (channels roughly correspond to intensity Y and two chrominance values, I and Q). We blur the intensity channel with a wide Gaussian filter. The resulting YIQ image is converted back to the RGB color space for output.

We noted above the tendency of many Retinex-based algorithms to operate exclusively by selectively brightening the image. This may lead an implementation to produce an over-enhanced image (Ng and Wang, 2011), which not only underscores the difficulty of applying Retinex (and the incredible capabilities of the human visual system to adapt) but also reinforces the notion that tuning parameters is a necessary condition for successful application of a Retinex-based algorithm. One solution to this problem has been to apply a post-processing step akin to gamma correction or a similar global operator to reduce the overall image brightness.

As with the Multi-scale Retinex algorithm, we find that some of the color quality of the resulting image may be lost. We apply a color restoration, but must also adapt it to the multi-center algorithm. In Multi-scale Retinex, color restoration always implied a darkening of the image, on the assumption that (since the algorithm would always brighten an image) the loss of color was due to an “over-shoot” in the brightening effect. Its color correction is of the form of Equation 4.

$$a_i(x, y) = \ln \left(1 + C \frac{I_i(x, y)}{\sum_{n=1}^N I_n(x, y)} \right)$$

Equation 4: Multi-scale Retinex color restoration

Instead of using the logarithmic ratio of each color channel with the sum of the channels, we scale by the ratio of each channel to the average of the N channels, using Equation 5. This ratio is then averaged with the value 1 to prevent overcorrection.

$$a_i(x, y) = \frac{1}{2} \left(1 + N \frac{I_i(x, y)}{\sum_{n=1}^N I_n(x, y)} \right)$$

Equation 5: Color restoration for adapted Retinex

RESULTS

We measure our results by applying measures of contrast, despite the questionable perceptual applicability of standard definitions of contrast (Peli, 1997). We began with the definition of local band-limited contrast (Peli, 1990) to measure the changes applied by our algorithm to images; we adapt this measure to compare differences between a pixel and its background (surrounding context). This differs from the ratio used in local band-limited contrast, which is typical of contrast calculations. We need a definition that allows us to compare cases in which the surround is darker than a pixel directly to cases in which the pixel is darker than its context.

Mean Local Band-limited Contrast of Target				
Image (Figure)	Orig	HE	MSR	AR
Landscape (5)	4.71	33.75	8.89	45.02
Skyline (6)	1.23	7.94	5.63	28.43
Snowy Road (7)	5.88	25.42	12.26	34.20
Sandstorm (9)	13.69	23.47	8.34	47.96
Underwater (10)	13.11	33.27	8.47	49.72
Night Vision (12)	32.68	44.46	17.46	58.34

Table 1: Numerical comparison of our algorithm to other approaches; Orig = metric computed on original image, HE = metric computed on image processed by histogram equalization, MSR = metric computed on image processed by Multi-scale Retinex algorithm, and AR = metric on image processed by our adapted Retinex processing.

We compare the metric results on the original image (Orig), the image produced by histogram equalization (HE), the image processed with the Multi-scale Retinex (MSR) algorithm (Rahman et al., 2004), and our adapted Retinex-based processing (AR). As seen in Table 1, our algorithm consistently increases the contrast measure to a greater extent than HE and MSR do.

The human visual system detects spatial differences in intensity; thus, one could argue from these numbers that there will be more information available in the image from adapted Retinex processing. How useful this information is for a particular task is a measurement that must be made either through subjective assessment or an objective task performance user study (Moroney and Tastl, 2004); finding an emergency suitable landing strip as a pilot could be such a task (Woodell et al., 2006).

We compute our comparison metric in a manually specified area of interest. We could specify the entire image; in one case below, this is precisely what we do. In some examples, we specify a portion of the scene as the area of interest, denoted by red rectangles in Figures 5-7. This introduces the only subjective aspect to our contrast-based comparison. Looking through the images, we get a sense of the subjective performance of the algorithm. On the Landscape image (Figure 5), AR performed notably better than MSR in removing the atmospheric haze, drawing out building details, unmasking the water tower near the upper left. Similarly, on the Skyline image (Figure 6), all three processing algorithms increase the visibility of the buildings in the target region, with AR resulting in greater noise than HE or MRS. In the Snowy Road scene (Figure 7), the detail on the back of the truck and the utility pole in the distance emerge only with AR processing.



Figure 5: Images of the Landscape scene; *Top left*: original image, *Top right*: histogram equalization, *Bottom left*: Multi-scale Retinex, *Bottom right*: adapted Retinex processing. The red rectangle denotes the area of interest, which in this case is the entire image. In the adapted Retinex image, note the success of haze removal, the level of detail on the buildings, and the appearance of the water tower near the upper left.




Figure 6: Images of the Skyline scene; *Top left*: original image, *Top right*: histogram equalization, *Bottom left*: Multi-scale Retinex, *Bottom right*: adapted Retinex processing. The red rectangle denotes the area of interest for objective evaluation. Note the greater detail in the region of interest, as well as the foreground. Adapted Retinex processing removed more haze, albeit with a perceptible increase of the noise in the resulting image.




Figure 7: Images of Snowy Road scene; *Top left*. original image, *Top right*. histogram equalization, *Bottom left*. Multi-scale Retinex, *Bottom right*. adapted Retinex processing. The red rectangle denotes the area of interest. Note that histogram equalization appears to reduce detail on the back of the truck. Adapted Retinex processing increases the contrast and makes the utility pole in the background much more salient.



Figure 8: A playground scene with good initial visibility has its contrast increased by adapted Retinex processing (right) without masking details visible in the original image (left).



Figure 9: Our algorithm cleans up a sandstorm from this image. The minor color banding artifacts (visible in the sky) seem to be a small price for the increase in visibility. The red rectangle denotes the area for objective evaluation.




Figure 10: Failure on underwater imagery: four white fish (two at top center, one at bottom center, and one at left center) are barely any more visible in the result (bottom) than in the original (top)

Another interesting test of AR processing is to see how it performs on an image that already has good visibility. Figure 8 (left) shows such an image of a playground scene. AR processing (right) results in brighter colors and a darkening of portions of the sky to increase the contrast between the clouds and the clear streaks of the sky. While the saturation of the colors is stark, the AR processing has not hidden any interesting details from the original image. Sandstorms provide another interesting test case on which AR processing performs well (Figure 9).

One unsuccessful test was on the class of underwater images. These images have scattering effects that may be mitigated by wavelet decomposition and estimation of optical properties of water (Hou et al., 2007). AR processing can maintain the color of some fish and correct the distorted film processing of the sunlight, but Figure 10 shows a failure on four white fish. Also, the noise in this image and other underwater images resulting from AR processing is too high for aesthetic purposes and perhaps for objective tasks as well. Noise filters designed for underwater images did not improve the overall results as either a pre-process or a post-process.



Figure 11: A successful example of adapted Retinex processing on a traditional brightening task



Figure 12: Applying the AR processing to night vision imagery enables the background, hidden in the original (left) to be seen after processing (right)

CONCLUSIONS

Our Retinex-inspired processing performs well on the task of brightening a dark image (Figure 11). An extreme example is seen in applying the algorithm to enhancing night-vision imagery (Figure 12). Note how the distant structures, invisible in the original, are clear in the processed image, albeit with an artifact introduced in the sky. The result also shows reduced over-exposure of the people.

We can identify avenues for further work. Success with adapted Retinex processing depends on the amount of noise in the input image. We have tested using median filters and adjusting the Gaussian filtering. A more powerful and general method of managing noise would further improve the algorithm's performance and reduce the need for pre- or post-processing. A completely automated

algorithm may not be necessary or even desirable for a particular task, but it still represents a worthy goal for general use of adapted Retinex processing.

As noted above, subjective testing in the context of a particular application or task is the true measure of the usefulness of the adapted Retinex processing. We make use of graphics hardware acceleration of the Fourier transform to increase the computational performance of our algorithm; further parallelization of the code could lead to real-time performance.

REFERENCES

Marcelo Bertalmío, Vicent Caselles, and Edoardo Provenzi, "Issues About Retinex Theory and Contrast Enhancement." *International Journal of Computer Vision* 83(1):101-119 (June 2009)

Brian Funt, Florian Ciurea, and John McCann, "Tuning Retinex Parameters." In Human Vision and Electronic Imaging VII, Proceedings of SPIE Vol. 4662 (January 2002)

Weilin Hou, Deric J. Gray, Alan D. Weidemann, Georges R. Fournier, and J.L. Forand, "Automated underwater image restoration and retrieval of related optical properties." In *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pgs. 1889-1892 (July 2007)

Edwin H. Land, "The Retinex Theory of Color Vision." *Scientific American* 237(6):108-128 (December 1977)

Edwin H. Land, "Recent Advances in Retinex Theory." *Vision Research* 26(1):7-21 (1986)

Edwin H. Land and John J. McCann, "Lightness and Retinex Theory" *Journal of the Optical Society of America* 61(1):1-11 (January 1971)

John J. McCann, "Capturing a black cat in shade: past and present of Retinex color appearance models." *Journal of Electronic Imaging* 13(1):36-47 (January 2004)

Nathan Moroney and Ingeborg Tastl, "Comparison of Retinex and iCAM for scene rendering." *Journal of Electronic Imaging* 13(1):139-145 (January 2004)

Michael K. Ng and Wei Wang, "A Total Variation Model for Retinex." *SIAM Journal on Imaging Sciences* 4(1):345-365 (March 2011)

Eli Peli, "Contrast in complex images." *Journal of the Optical Society of America A* 7(10):2032-2040 (October 1990)

Eli Peli, "In Search of a Contrast Metric: Matching Perceived Contrast of Gabor Patches at Different

Phases and Bandwidths." *Vision Research* 37(23):3217-3224 (December 1997)

Zia-ur Rahman, Daniel J. Jobson, and Glenn Woodell. "Retinex processing for automatic image enhancement." *Journal of Electronic Imaging* 13(1):100-110 (January 2004)

Glenn A. Woodell, Daniel J. Jobson, Zia-ur Rahman, and Glenn D. Hines, "Advanced Image Processing of Aerial Imagery." In Visual Information Processing XV, Proceedings of SPIE Vol. 6246 (April 2006)

ACKNOWLEDGMENTS

The authors wish to thank Jonathan Decker, Peter Gabor, Irwin E. Sobel, Jack Bono, and Leslie Smith. This work was supported in part by the NRL Base Program and in part by the ONR/ASEE Science and Engineering Apprenticeship Program.

The views expressed herein are the personal opinions of the authors and are not necessarily the official views of the Department of Defense or any military department thereof.

Mark A. Livingston, Ph.D., is a computer scientist at the Naval Research Laboratory in the Information Management and Decision Architectures Branch. He manages and conducts research in interactive computer graphics, including virtual environments, augmented reality, image processing, visual analytics, and scientific / information visualization. He received his doctorate from the University of North Carolina at Chapel Hill in 1998.

Caelan R. Garrett participated in the Science and Engineering Apprenticeship Program at NRL in 2010. An early version of this work was recognized in the national finals of the Siemens Competition in Math, Science, and Technology. Mr. Garrett graduated Thomas Jefferson High School for Science and Technology in 2011 and is currently a freshman at Massachusetts Institute of Technology.

Zhuming Ai, Ph.D., is a computer engineer at the Naval Research Laboratory in the Information Management and Decision Architectures Branch. He conducts research on virtual environments, augmented reality, user interfaces, computer graphics, and image processing. Before joining NRL, he was a research assistant professor at the Univ. of Illinois-Chicago, doing research on virtual reality in medicine. He received his doctorate from Southeast University, Nanjing, China.