Feature visibility limits in the non-linear enhancement of turbid images

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

[‡]NASA Langley Research Center, Hampton, Virginia 23681.
[†]College of William & Mary, Department of Computer Science, Williamsburg, VA 23187.

ABSTRACT

The advancement of non-linear processing methods for generic automatic clarification of turbid imagery has led us from extensions of entirely passive multiscale Retinex processing to a new framework of active measurement and control of the enhancement process called the Visual Servo. In the process of testing this new non-linear computational scheme, we have identified that feature visibility limits in the post-enhancement image now simplify to a single signal-to-noise figure of merit: a feature is visible if the feature-background signal difference is greater than the RMS noise level. In other words, a signal-to-noise limit of approximately unity constitutes a lower limit on feature visibility.

1. INTRODUCTION

In the previous development of non-linear image enhancement methods,[1–3] our goal was to enhance the visual realism of the recorded digital image to more closely approach the generally much better visibility of direct scene perception by the human observer. For images acquired under turbid—fog, smoke, haze, snow, rain—imaging conditions, there is already a close parity between the recorded image and the direct observation. So the goal of enhancement now becomes fundamentally different: we wish to greatly exceed the performance of the human observer. This is of particular interest for enhancing imagery acquired under turbid aviation conditions. A generic automatic computation that does this provides the enabling technology for real-time image enhancement that can be projected to the pilot's heads-up display (HUD). This type of imagery is especially important in commercial aviation during runway approach and landing and in general aviation during the entire flight sequence from take-off to landing.

Our interest in enhancing images acquired under turbid imaging conditions, coupled with the scientific insights[4] gained from previous purely passive retinex processing led us to formulate a more comprehensive framework of active measurement and control of the image enhancement process: the Visual Servo (VS). The major lesson learned from these scientific insights was that the good visual representations produced by retinex processing all converged uniquely to an ideal statistical characterization. This, together with additional constraints, led to the formulation of an entirely new set of visual measures for image contrast, lightness and sharpness. These measures form the basis for VS controls that affect the level of image enhancement and, hence, image quality. The VS additionally contains a special module for detecting turbid imaging conditions and invokes special processing to produce maximal scene feature clarity.

This framework represents a form of visual intelligence: the software quantitatively assesses visual quality before and after each enhancement step, and guided by these measurements, strives to achieve a standard high level of visual quality. The VS controls still rely on non-linear image processing elements, so conventional end-to-end systems analyses[5–8] cannot be employed to characterize the imaging and computing scheme as a whole. Therefore we seek to understand how to characterize performance in lieu of having linear systems analysis as a tool. In this paper we describe the Visual Servo concept, present results for diverse turbid imaging conditions to indicate its generic performance as an automatic computation, and examine the critical issue of defining a figure of merit for the post-enhancement feature visibility limit in turbid imaging conditions.

Contact: DJJ: d.j.jobson@larc.nasa.gov; ZR: zrahman@as.wm.edu; GAW: g.a.woodell@larc.nasa.gov;

2. THE VISUAL SERVO (VS) CONCEPT

Our extensive previous experience with retinex image enhancement,[3, 9] led us to conclude that further improvements in enhancement were possible in terms of contrast and image sharpness and that the extreme narrow dynamic range case of turbid imaging could be also be encompassed by a fundamental shift in approach away from purely passive retinex processing to an active measurement and control system. The scientific insights[4] gained from the experience of retinex processing of very large numbers of highly diverse images provided the foundations for constructing entirely new visual measures for image contrast, lightness, and sharpness. Underlying this effort is the core idea that there is an ideal visual representation with consistent statistics for recorded images[3, 4] and that the enhancement process is one of trying to make any image approach this ideal as closely as possible. This core idea also relates to the visual inadequacy of the linear representation of recorded image data. The idea rejects the notion of imaging as a replication process with quality defined by minimizing artifacts, and shifts to the notion of imaging being a (highly non-linear) transformation process that seeks to achieve a good visual representation whose statistics depart sharply from those of the linear representation.

With this background, the study of the overall global statistics of good visual representations revealed that global statistics alone could not support the definition of visual measures.[4] There was insufficient capture of the visual sense of contrast and lightness. However, regional spatial ensembles did provide a basis for quantifying visual contrast and lightness. These were augmented by the development of a new sharpness measure which together capture the most comprehensive and key visual elements of images. The measures have been extensively tested, but are not yet ready for full technical exposition, so here we will discuss them and the resulting VS at the conceptual and schematic level.

The Visual Servo is shown in Figure 1. The reason for calling the computation a "servo" stems from the fact that it is based upon similar ideas to electro-mechanical servo systems of active measurement and control. The basic flow for enhancing an image is as follows:

1. measure a key visual parameter

2. based upon the measured value, activate an enhancement control to improve the overall brightness, contrast and sharpness of the image

3. recompute the visual measure

4. if the measured value of the parameter achieves the high visual standard then terminate the process

5. otherwise, iterate until either the visual standard has been achieved or the VS has determined internally that all reasonable enhancement processing has been exhausted
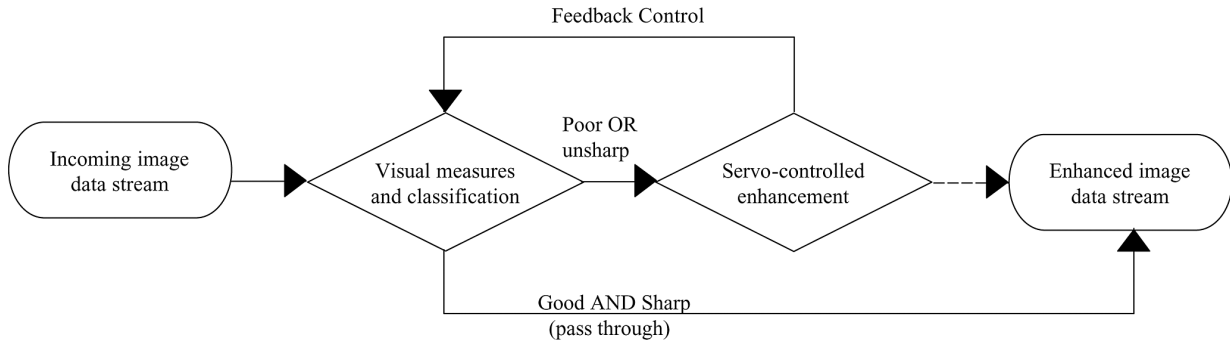
Figure 1: The automatic visual servo.

Images acquired under turbid imaging conditions represent a special case for detection and processing due to their extremely narrow, but unpredictable, dynamic range. For these images, the VS determines when this

very low contrast is occurring, and then invokes custom processing to achieve a very powerful enhancement. This enhancement, from the results to be shown in Figure 2, appears to produce maximum feature contrast that is limited only by the noise inherent to the imaging process.

3. VISUAL SERVO RESULTS FOR TURBID IMAGING

Turbid imaging covers a very wide range of imaging conditions where there is obscuration between the scene and the imaging sensor due to particle scattering in the imaging medium. For atmospheric turbidity, the veiling can be due to fog, smoke, haze, dust, rain, or snow, or some combination of these, such as smog. For underwater imaging the turbidity is most often due to suspended cellular plant life, sediment particles, or some combination thereof. In order to validate the generality of the computation for turbid imaging conditions, we have tested the VS on images with very varied types and degrees of turbidity, as well as highly diverse scene content. The result is an automatic, general purpose computation that is as applicable to aerial imagery as it is to underwater imagery. All of the results shown in Figure 2 (right column) were obtained in the fully automatic mode, without any additional tweaking of parameters or post-processing. Figure 2 (middle column) shows the performance of the default MSRCR on the same images. It is evident from comparing the two columns why we moved away from the passive MSRCR to the active VS for enhancing images acquired in turbid imaging conditions.

The turbidity detection and custom processing of the VS has worked extremely well in (almost) all the cases that we have tested thus far. For lightly turbid conditions, the VS transitions smoothly to the non-turbid enhancement processing, invoking different servo modules that provide "weaker" enhancements. The VS has been tested with many hundreds of still color images, as well as with color and long, and short wave infrared (IR) video imagery. The IR imagery was acquired from aviation sensors during test flights, or from sensors mounted on a 250 feet high gantry to simulate flight conditions such as long throw views of landing approaches.

The performance of the VS has been outstanding for still and video imagery.* Additionally, it does not seem to be sensitive to the type of particulate scattering involved. Good clarity—visibility distance improvement— was achieved for moderate fog, severe haze, moderately thick smokes, heavy rain and snow as well as for moderately thick underwater turbidities. Some clarity was possible for heavy fog conditions (see Figure 4). This latter performance limitation in dense fog was true for all imagery types tested—color, short wave IR, and long wave IR. For all but the dense fog case, sufficient clarity was achieved so that often all traces of obscuration were removed. For cases of thicker turbidity, visibility as an increase in feature distance visibility was greatly improved. For the cases where we acquired the image data ourselves, we could compare the performance of the VS to what we had observed at the time of acquisition. In all cases, except severe fog, the VS result was far better than our observed visibility.

A reasonable baseline for performance comparison is to compare servo results with those for the conventional automatic histogram modification method—autolevels. Autolevels is a moderately powerful automatic image enhancement technique which is quite useful for images with low contrast that do not contain regions of saturation. It is a histogram stretch technique that adjusts the dynamic range of the displayed image based upon a fixed parameter that determines the significant dynamic range of the input image. The dynamic range compression adjustment of the Autolevels process is very different from the intrinsic dynamic range compression that the non-linear processing embedded in the VS provides. So the primary performance differences to be expected are the result of the non-linear dynamic range compression. Figure 3 shows examples of a comparison between the performance of Autolevels and the VS. It is clear from the figure that the VS has much better performance than Autolevels. This implies that even for very narrow dynamic range imaging, the non-linear dynamic range compression is still quite advantageous. The reason for this is that the very narrow dynamic range is not stable regionally across the image. There is still lot of "shading" variation due to spatial lighting effects within the narrow dynamic range of the turbid image. Of course there can be cooperative cases where the lighting happens to be spatially uniform, and for these cases, the Autolevels performance will approach that of the VS. The servo performance though should be much more all encompassing of the full complexity of real turbid imaging conditions where shadows, and other lighting variations, as well as highly spatially variable degrees of turbidity are going to be encountered within a particular image frame.
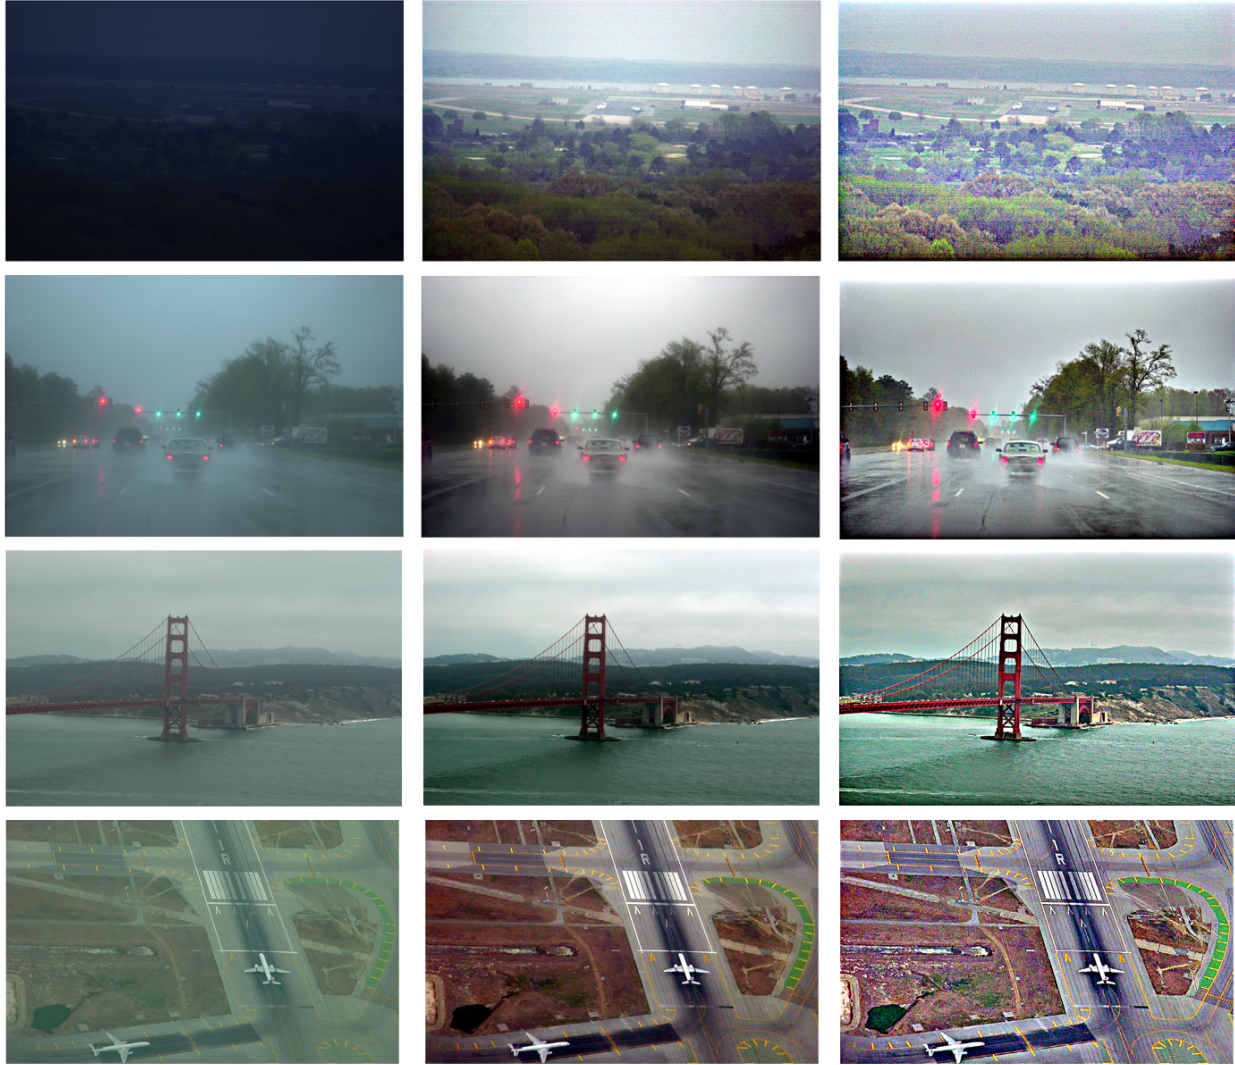

*Further examples of VS enhancements can be found on `http://dragon.larc.nasa.gov/retinex`.
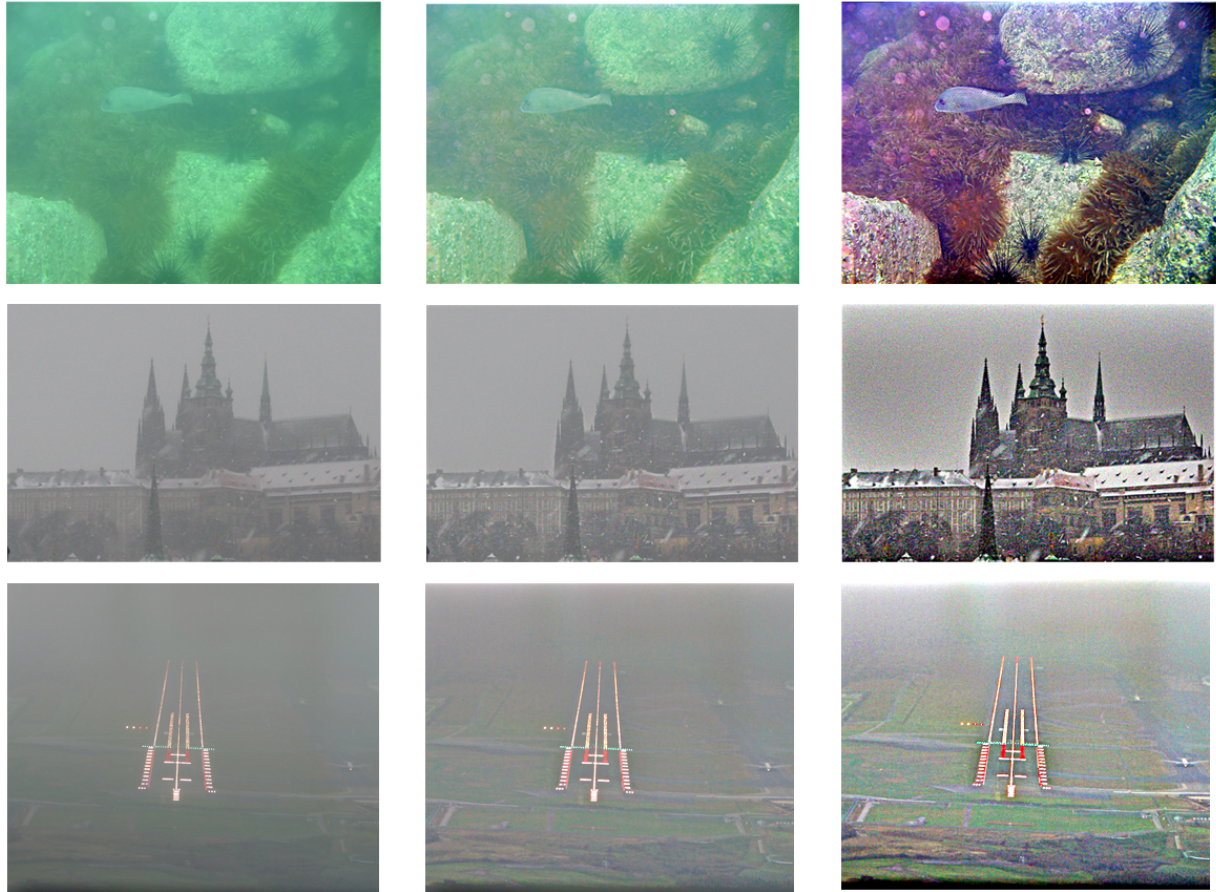
Figure 2. The performance of the VS and the default MSRCR on images acquired under turbid imaging conditions. The left column shows the original images; the middle column the default MSRCR processed images; and the right column, images that have been processed using the VS.

Figure 3. A comparison of automatic visual servo with autolevels. The left column shows the original images; the center column shows the Autolevels enhanced images; and the right column shows the enhancements produced by the Visual Servo.
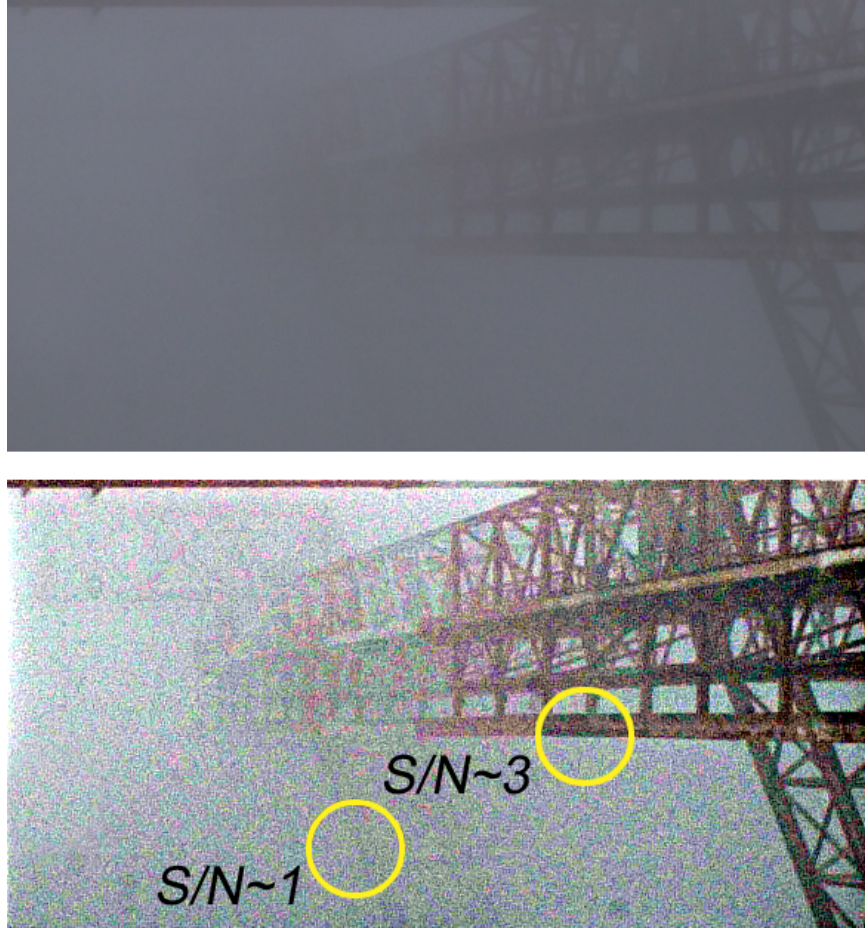
Figure 4: Feature visibility limit of the visual servo

4. THE POST-ENHANCEMENT FEATURE VISIBILITY LIMIT

The dense fog case provides an instructive example for defining feature visibility limits since we encounter post-enhancement opacity immediately and can readily track feature visibility deterioration within short distances. This is shown in Figure 4 where the visibility deterioration occurs for a feature signal-difference approximately equal to the root-mean-square (RMS) noise in the post enhancement image. For slightly closer distances or slightly less turbidity, feature visibility improves rapidly. Even for very low RMS feature signal-difference-to-noise ratios (≈ 3, the visibility is remarkably good as shown in Figure 4). This result appears consistently in a number of other highly turbid test images, so we conclude that post-enhancement feature visibility can be defined by this very simple figure of merit. Unlike more complex figures of merit which must account for both feature signal level as well as feature signal difference, the post enhancement image domain is governed by this noise limit alone. For sensors with lower noise, there will be a consequent improvement in visibility, and ultimately the visibility limit should be set by the signal photon noise, or other scene noises (such as random variations in scattering particle densities) for extremely high sensitivity sensors. This latter case assumes that digitization is done at sufficiently high bit levels so that quantization noise is made lower than any scene noise sources.

5. CONCLUSIONS

An outgrowth of previous developments in non-linear image enhancement was the realization that further improvements in visual contrast and sharpness were needed and that the case of turbid imaging needed to be addressed especially as a significant case for aviation imaging. For this case, a computation was needed which is automatic so that it can be implemented in real-time hardware for enhanced pilot vision during poor visibility flight conditions. These issues together with new scientific insights gained from retinex image processing, led us to develop the VS which is more comprehensive than the previous passive retinex image processing. A set of measures of visual contrast, lightness and sharpness were defined and tested which serve as the basis of the active measurement and control VS system. The system does use non-linear image enhancement for the servo control modules and therefore does not have performance that is derivable from the usual linear systems analysis. Therefore we have extensively tested the VS of wide ranging types and degrees of image turbidity as well as in general purpose imaging.

The servo computation performs well in all turbid imaging condition short of dense fog, and greatly exceeds the human observer's direct perception in all but the dense fog case. The servo handles all manner of moderate fogs, severe hazes, heavy rain, smoke, and heavy snow conditions quite well. It has been tested on color still images as well as color and FLIR video from aviation sensors in varied flight conditions—in and out of clouds, severe haze, twilight-haze, moderate fogs and smoke among others.

Given the use of non-linear image processing, a major issue is how to define a figure of merit for feature visibility limits. An experimental study of post-enhancement image data revealed that feature visibility was limited solely by a very simple figure of merit compared to those used for unenhanced imagery. This figure of merit is that features are visible post-enhancement as long as the feature/background signal-to-noise ratio is greater than unity. Feature visibility increases rapidly for higher signal-to-noise ratios such that a S/N of ≈ 3 has quite good visibility for example.

While our primary interest for the VS is clarifying aviation imagery during poor visibility flight conditions, the servo performs well on underwater turbid images, and in poor visibility road conditions during driving. Therefore we expect that computation has applications to a variety of turbid imaging conditions where a human observer needs to be augmented visually to improve safety and visual performance.

Acknowledgments

REFERENCES

1. 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.
2. 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.
3. Z. Rahman, D. Jobson, and G. A. Woodell, "Retinex processing for automatic image enhancement," in *Human Vision and Electronic Imaging VII*, B. E. Rogowitz and T. N. Pappas, eds., pp. 390–401, Proc. SPIE 4662, 2002.
4. D. 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.
5. F. O. Huck, C. L. Fales, R. E. Davis, and R. Alter-Gartenberg, "Visual communication with Retinex processing," *Applied Optics* , Apr 2000.
6. 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.

7. C. L. Fales, F. O. Huck, R. Alter-Gartenberg, and Z. Rahman, "Image gathering and digital restoration," *Philosophical Transactions of the Royal Society of London A* **354**, pp. 2249–2287, Oct. 1996.
8. J. A. McCormick, R. Alter-Gartenberg, and F. O. Huck, "Image gathering and restoration: Information and visual quality," *Journal of the Optical Society of America A* **6**(7), pp. 987–1005, 1989.
9. 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.