A Spatial Post-Processing Algorithm for Images of Night Scenes

William B. Thompson and Peter Shirley
University of Utah

James A. Ferwerda
Cornell University

Abstract. The standard technique for making images viewed at daytime lighting levels look like images of night scenes is to use a low overall contrast, low overall brightnesses, desaturation, and to give the image a "blue shift". This paper introduces two other important effects associated with viewing real night scenes: visible noise, and the loss of acuity with little corresponding perceived blur.

1 Introduction

While many researchers have studied how to display images of daylit (photopic) scenes (a good overview can be found in (Durand and Dorsey 2000)), little work has dealt with the display of images of night (scotopic) scenes. The work that has been done uses some combination of a reduction in contrast/brightness, desaturation, a blue-shift, and a low-pass filter (Upstill 1995; Tumblin and Rushmeier 1993; Ward 1994; Ferwerda, Pattanaik, Shirley, and Greenberg 1996). These strategies have also been used for many years in film (Samuelson 1984).

There are two subjective properties of night viewing that are missing from previous techniques. First is that although there is a loss-of-detail at night, there is not a sense of blurriness. Second is the somewhat noisy appearance of night scenes. Both of these effects have only been partially explored scientifically (Makous 1990; Field and Brady 1997).

In this paper we present techniques for augmenting images with the loss-of-detail and noise effects associated with night vision. Because the the subjective properties of scotopic effects are themselves not fully understood, and because we have the added confound of trying to mimic these effects in photopically-viewed images, we will not use a specific quantitative model to drive our method. Rather, we develop an empirical image-processing technique that is consistent with what is known, and highlight the the features of the techniques by demonstrating them on a test pattern and real images.

2 Background

While scotopic vision has not been studied to the extent that photopic vision has, there is still a wealth of knowledge about it (e.g., see (Hess, Sharpe, and Nordby 1990)). We review some of that work here, and conclude that while it does not provide us with an explicit model to control our image processing, it does give us subjective information that is useful.

2.1 Visual acuity and perception of blur

Visual acuity is reduced in scotopic viewing due in part to the relative sparcity of rods vs. cones in the central retina. Subjectively, however, night does not appear blurry (Hess 1990). This has been observed in people with only rod vision ("complete achromats") as well; Knut Nordby, a perceptual psychologist with that condition writes:

> I experience a visual world where things appear to me to be well-focused, have sharp and clearly defined boundaries and are not fuzzy or cloudy. (Nordby 1990)

Simulating such loss of acuity by low-pass filtering the original image (e.g., (Ferwerda, Pattanaik, Shirley, and Greenberg 1996)) fails to effectively capture the appearance of night scenes because when viewed in bright light, the images appear blurred rather than having the clearly defined boundaries Nordby describes.

Surprisingly little is know about the image cues that generate a sense of blur in the human vision system. While the power spectrum of natural images falls with frequency at about f^{-2}(e.g., (Ruderman 1997)), there is substantial variability from image to image, even when the images appear sharp (Field and Brady 1997). As a result, a simple frequency distribution analysis will fail to predict the appearance of blur. Field and Brady argue that a person's sensitivity to blur depends on both the high frequency content of an image and the density of fine-scale edges. In particular, an image will tend not to look blurred if fine detail edges are sufficiently crisp, independent of the number of fine-scale edges in the image.

2.2 Noise

Night scenes also appear somewhat noisy subjectively. This is not surprising if one thinks of the scotopic vision system as a gain-control amplifier where low signals will have a noisy output (Lamb 1990). However, there are more potential sources of noise than that simple explanation suggests. These include the quantum number of photons, noise originating in the receptors themselves, and noise in the neural circuitry (Sharpe 1990). These various types of noise are themselves only partially quantified, and they are composed in complex ways which is not fully understood (Makous 1990). The noise is additive, but can behave as if multiplicative because neural noise might be added after logarithmic gain control (Makous 1990).

3 Night filtering

One way to simulate the loss of acuity in scotopic viewing would be to use some form of anisotropic diffusion (e.g., (Tumblin and Turk 1999)), which has the effect of smoothing the geometry of edges without blurring across the edges. We have adopted a spatial filtering approach as an alternative because it does not require an iterative process and because the frequencies preserved can be more easily controlled. As noted above, an image will tend not to look blurred if fine detail edges are sufficiently crisp, independent of the number of fine-scale edges in the image. For an original image I scaled to the range $[0.0 - 1.0]$, we start by low pass filtering the image, using convolution with a Gaussian kernel G_{blur}, where the standard deviation σ_{blur} is chosen to remove fine-scale detail that would not be visible at night:

$$I_{\text{blur}} = G_{\text{blur}} * I \tag{1}$$

It is not sufficient to simply apply a standard sharpening operator such as a narrow support unsharp mask to this blurred image, since there are no longer high frequencies to emphasize. A broader band sharpening filter, such as a larger extent unsharp mask, produces noticeable ringing. Instead, we apply a sharpening operator tuned to the finest detail edges remaining after the initial low pass filtering operation. This is done by first creating a bandpass filtered image using the difference-of-Gaussian method, where the filter selects out the highest remaining spatial frequencies. The original image is convolved with a second Gaussian kernel G_{blur_2}, with standard deviations $\sigma_{\text{blur}_2} = 1.6\ \sigma_{\text{blur}}$, after which we take the difference of the two blurred images:

$$\begin{aligned} I_{\text{blur}_2} &= G_{\text{blur}_2} * I && (2) \\ I_{\text{diff}} &= I_{\text{blur}} - I_{\text{blur}_2} && (3) \end{aligned}$$

I_{diff} is a close approximation to the $\nabla^2 G$ function used in some edge detectors and is a bandpassed version of I (Marr and Hildreth 1980). I_{blur} can be decomposed into the bandpassed component plus an even lower-pass component:

$$I_{\text{blur}} = I_{\text{blur}_2} + I_{\text{diff}} \tag{4}$$

The appearance of blur can be substantially reduced while still attenuating fine detail by sharpening the edges in the bandpassed component. One way to do this would be to multiply I_{diff} by an appropriately chosen constant $\alpha > 1.0$. This is similar to a large-neighborhood unsharp mask and produces the same problems with ringing. Instead, we exploit the fact that edges are located at zero crossings of I_{diff} (Marr and Hildreth 1980). Edge sharpening can thus be accomplished by increasing the contrast of I_{diff} for values near 0:

$$I_{\text{night}} = I_{\text{blur}_2} + {I_{\text{diff}}}^{1.0/\gamma_{\text{edge}}} \quad \text{for} \quad \gamma_{\text{edge}} > 1.0 \tag{5}$$

The value of γ_{edge} affects the apparent crispness of the final edges. A value of 1.25 works well for a wide range of naturally occurring images.

The literature on dark noise is not sufficiently definitive as to allow the development of a precise physiological or psychophysically based approach to transforming images so that they look more night-like, even when viewed in bright light. The problem is further compounded by the fact that different types of display devices can have very different effects on the appearance of high frequency image noise. As a result, we have adopted an approach in which we add zero-mean, uncorrelated, Gaussian noise to images after they have been spatially filtered, with the standard deviation, σ_{noise} of the noise adjusted subjectively depending on the display device. ($\sigma_{\text{noise}} = .0125$ for the examples presented in this paper.)

4 Implementation issues

We assume the input is a displayable RGB image. If our source image is a high-dynamic range image, it should first be tone-mapped using one of the standard techniques (e.g., (Reinhard, Stark, Shirley, and Ferwerda 2002)). This RGB image is then mapped to a scotopic luminance image, where each pixel is a single number indicating the "brightness" of the pixel as seen at night. This will tend to favor blues because the rods are more sensitive to blues than to greens and reds. This can either be done heuristically using a linear weighted sum of the the RGB channels or can be done by first converting to XYZ using one of the standard manipulations such as:

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 0.5149 & 0.3244 & 0.1607 \\ 0.2654 & 0.6704 & 0.0642 \\ 0.0248 & 0.1248 & 0.8504 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \tag{6}$$

We can then use an empirical formula to approximate the scotopic luminace V (Larson, Rushmeier, and Piatko 1997):

$$V = Y \left[1.33 \left(1 + \frac{Y + Z}{X} \right) - 1.68 \right]. \tag{7}$$

This results in a single-channel image that ranges to values in approximately the [0,4] range. That image should then be scaled and multiplied by a bluish grey to give an unfiltered "night image". Some authors suggest a blue with chromaticity approximately (0.03,0.03) below the white point (Durand and Dorsey 2000), although many films have a more saturated blue in their day-for-night scenes. Thus for each pixel we have:

$$c_{\text{night}} = kVc_{\text{blue}} \tag{8}$$

where k and c_{blue} are chosen empirically.

The resulting bluish image is then spatially processed using the methods discussed in Section 3. If some of the pixels are above the scotopic level, e.g., a car headlight or an area near a flame, then we can combine a day and a night image. If we consider the original RGB image to have a color c_{day} at each pixel, then we can use a scotopic fraction s at each pixel and blend the images pixel by pixel:

$$c = sc_{\text{night}} + (1 - s)c_{\text{day}} \tag{9}$$

The fraction s can be chosen based on V, or can be chosen in a more principled manner (e.g., (Ferwerda, Pattanaik, Shirley, and Greenberg 1996)).

5 Examples

We assume that the images we process have already undergone some day-for-night processing (reductions in brightness, contrast, and saturation, together with a blue shift). The images in this section were first transformed using a photo editing tool rather than a more formal tone-mapping algorithm. However, the techniques should be applicable to any initial day-for-night image.

Figure 1 illustrates the effect of applying the spatial filtering portion of the method to a test pattern consisting of gratings of different scales. σ_{blur} was chosen such that the individual bars in the center grouping would be resolvable, while the fine detail of the right-most bars would not. Figure 2 shows intensity plots of the test grating. The difference between night filtering and simple blurring is most apparent when comparing the mid-sized bars in Figures 2d and 2e. The night filtering does more than enhance the contrast: the definition of the bars relative to the overall contrast is significantly increased, which is why Figure 1d looks crisper than Figure 1c.

Figures 3–8 illustrate the method on a natural image. Prior to spatial processing, the original image was corrected to account for photographic and digitization non-linearities, with γ set based on a calibration target. Figure 4 was produced from Figure 3 by a manual day-for-night mapping, involving a reduction in brightness, contrast, and saturation, plus with a blue shift. Figure 5 shows what would happen if the loss of acuity associated with viewing the scene at night was simulated by blurring. Figure 6 preserves the same level of detail as Figure 5, without appearing to be nearly as blurry. Figure 8 shows the effects of adding noise to Figure 6.

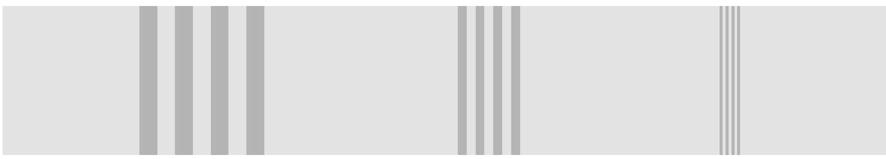

(a) Original image

(b) Day-for-night tone mapping

(c) Blurred to remove fine detail

(d) Night filtered to preserve same level of fine detail

Figure 1: Test grating: The same level of detail is resolvable with Gaussian blurring and night filtering, but the night filtered image looks sharper

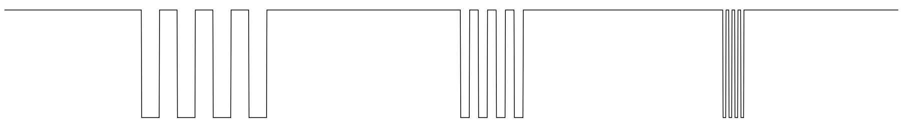

(a) Plot of original intensities

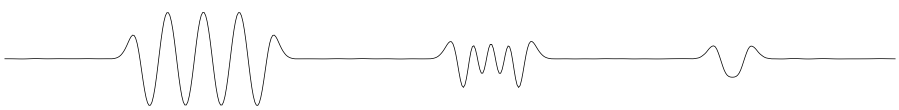

(b) Plot of $I_{\text{diff}} = I_{\text{blur}} - I_{\text{blur}_2}$

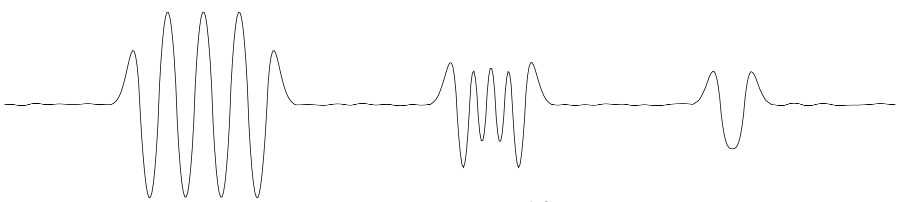

(c) Plot of $I_{\text{diff}}{}^{\frac{1.0}{\gamma}}$

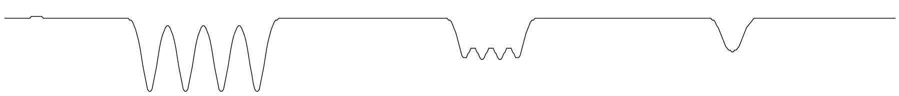

(d) Plot of blurred image

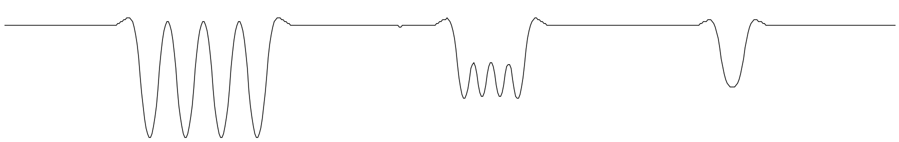

(e) Plot of night filtered image preserving same level of detail

Figure 2: Plots of test grating images

Figure 3: Original image

Figure 4: Day-for-night tone mapping

Figure 5: Blurred to remove fine detail

Figure 6: Night filtered, with the same level of fine detail as in Figure 5

Figure 7: Blurred plus noise

Figure 8: Night filtering plus noise

6 Discussion

There are three parameters of the operator: σ_{blur} which is physically based (smallest resolvable detail), and γ_{edge} and σ_{noise} that are set subjectively. Experience suggests that a single setting is effective over a wide range of differing imagery. Because edges are stable over time in an image sequence, the edge operator we use will preserve locality and will also be stable (Marr and Hildreth 1980). Thus the operator is well-suited for animations. While the operator is designed for non-scientific applications, it might be calibrated to be useful in preference to a linear blur for low-vision simulations.

References

Durand, F. and J. Dorsey (2000). Interactive tone mapping. In *Eurographics Workshop on Rendering*, pp. 219–230.

Ferwerda, J. A., S. Pattanaik, P. S. Shirley, and D. P. Greenberg (1996). A model of visual adaptation for realistic image synthesis. In *Proceedings of SIGGRAPH*, pp. 249–258.

Field, D. J. and N. Brady (1997). Visual sensitivity, blur and the sources of variability in the amplitude spectra of natural scenes. *Vision Research 37*(23), 3367–3383.

Hess, R. F. (1990). Rod-mediated vision: role of post-receptorial filters. In R. F. Hess, L. T. Sharpe, and K. Nordby (Eds.), *Night Vision*, pp. 3–48. Cambridge: Cambridge University Press.

Hess, R. F., L. T. Sharpe, and K. Nordby (Eds.) (1990). *Night Vision*. Cambridge: Cambridge University Press.

Lamb, T. D. (1990). Dark adaptation: a re-examination. In R. F. Hess, L. T. Sharpe, and K. Nordby (Eds.), *Night Vision*, pp. 177–222. Cambridge: Cambridge University Press.

Larson, G. W., H. Rushmeier, and C. Piatko (1997, October - December). A visibility matching tone reproduction operator for high dyna mic range scenes. *IEEE Transactions on Visualization and Computer Graphics 3*(4), 291–306.

Makous, W. (1990). Absolute sensitivity. In R. F. Hess, L. T. Sharpe, and K. Nordby (Eds.), *Night Vision*, pp. 146–176. Cambridge: Cambridge University Press.

Marr, D. and E. Hildreth (1980). Theory of edge detection. *Proc. Royal Society London B 207*, 187–217.

Nordby, K. (1990). Vision in the complete achromat: a personal account. In R. F. Hess, L. T. Sharpe, and K. Nordby (Eds.), *Night Vision*, pp. 3–48. Cambridge: Cambridge University Press.

Reinhard, E., M. Stark, P. Shirley, and J. Ferwerda (2002). Photographic tone reproduction for digital images. In *Proceedings of SIGGRAPH*.

Ruderman, D. L. (1997). The statistics of natural images. *Network: Computation in Neural Systems 5*(4), 517–548.

Samuelson, D. W. (1984). *Motion Picture Camera Techniques* (second ed.). Focal Press.

Sharpe, L. T. (1990). The light adaptation of the human rod vision system. In R. F. Hess, L. T. Sharpe, and K. Nordby (Eds.), *Night Vision*, pp. 49–124. Cambridge: Cambridge University Press.

Tumblin, J. and H. Rushmeier (1993, November). Tone reproduction for computer generated images. *IEEE Computer Graphics and Applications 13*(6), 42–48.

Tumblin, J. and G. Turk (1999). LCIS: A boundary hierarchy for detail-preserving contrast reduction. In *Proceedings of SIGGRAPH*, pp. 83–90.

Upstill, S. (1995, August). *The Realistic Presentation of Synthetic Images: Image Processing in Computer Graphics*. Ph. D. thesis, University of California at Berkeley.

Ward, G. (1994). A contrast-based scalefactor for luminance display. In P. Heckbert (Ed.), *Graphics Gems IV*, pp. 415–421. Boston: Academic Press.

Acknowledgement

Henrik Jensen, Daniel Kersten, Gordon Legge, Jack Loomis, and Simon Premoze provided helpful comments and discussion. This work was supported in part by National Science Foundation grants 89-20219, 95-23483, 98-18344 and 00-80999.