Visual Acuity in Day for Night

Gloria Haro , Marcelo Bertalmío and Vicent Caselles

({gloria.haro,marcelo.bertalmio,vicent.caselles} @upf.edu)

Departament de Tecnologia. Universitat Pompeu Fabra (Barcelona, Spain)

Abstract.

In film production, it is sometimes not convenient or directly impossible to shoot some night scenes at night. The film budget, schedule or location may not allow it. In these cases, the scenes are shot at daytime, and the 'night look' is achieved by placing a blue filter in front of the lens and under-exposing the film. This technique, that the American film industry has used for many decades, is called 'Day for Night' (or 'American Night' in Europe.) But the images thus obtained don't usually look realistic: they tend to be too bluish, and the objects' brightness seems unnatural for night-light. In this article we introduce a digital Day for Night algorithm that achieves very realistic results. We use a set of very simple equations, based on real physical data and visual perception experimental data. To simulate the loss of visual acuity we introduce a novel diffusion Partial Differential Equation (PDE) which takes luminance into account and respects contrast, produces no ringing, is stable, very easy to implement and fast. The user only provides the original day image and the desired level of darkness of the result. The whole process from original day image to final night image is implemented in a few seconds, computations being mostly local.

Keywords: Day for night, visual perception, dark adaptation, Weber's Law, non-linear diffusion

1. Introduction

In his book 'Making movies' (Lumet, 1995), the great American director Sidney Lumet points out the difficulties of night shooting on location. If the geography of the location is very inconvenient or dangerous for shooting at night, or if the budget can't afford extra pay for night shooting, or if a schedule delay is out of the question, what is usually done is Day for Night: the night scene is shot at day, but with a technique that gives a 'night look' to the film.

This technique is exclusively optical, typically a blue filter is placed behind the lens, and the film is under-exposed. While this usually works, the results nonetheless lack realism.

Blue turns out to be the predominant color, the other colors virtually indistinguishable. The objects in the scene have a very unnatural brightness. Moreover, no detail is lost: if there is a sign with small fonts, we can read it as if it were day.

The problem is that with mere optical means (blue filters, under-exposure) we cannot expect to reproduce all the modifications that account for how we see the world at night. The human visual system works very differently under low light conditions. The colors are perceived as less vivid, brightness is modified according to wavelength, the contrast changes significantly, and visual acuity decreases. It is therefore usual for Day for Night footage to be heavily retouched in post-production, mostly manually, by a color artist.

In this article we propose an algorithm to automatically transform a 'day image' into a 'night' version of it. The only other input necessary is the desired level of darkness of the final result. Our algorithm uses a very simple set of equations to model the different factors involved in night vision. These equations are based on physical data and visual perception experimental data.

The main contributions of this paper are the following. We start by replacing daylight illumination with night-like illumination, a procedure which to the best of our knowledge is novel when addressing the Day for Night problem. For the simulation of the loss of visual acuity, we introduce a novel diffusion PDE that models the spatial summation principle (Cornsweet and Yellott, 1985), takes luminance into account and respects contrast, produces no ringing, is stable, very easy to implement and fast.

2. The Day for Night algorithm

Our algorithm takes as input a color RGB image (coming from digital video or obtained from a scanned film) and a desired level of luminance for the final result. We transform the image in five steps.

We perform the following operations for each pixel in the image, one at a time. In the first step, we estimate the reflectance values for the object in the scene at that particular pixel, so we can replace the daylight illuminant used to shoot the scene with a nighttime illuminant of choice. In the second step we modify the chromaticity of the estimated reflectance values, assuming that the eye is dark adapted. In the third step we modify the luminance values. In the fourth step we modify the contrast, since threshold values for 'just noticeable differences' depend on illumination levels. These steps 2 to 4 are implementations of algorithms already introduced in the literature on tone reproduction and modeling of visual perception, like those of (Tumblin and Rushmeier, 1993; Ward et al., 1997; Ferwerda et al., 1996). Finally, in the fifth step we perform diffusion on the image to account for the loss of visual acuity at nighttime illumination levels. For this last step we introduce a novel equation that models the spatial summation principle (Cornsweet and Yellott, 1985), but without producing ringing or other visual artifacts.

One note concerning the evaluation of the results. For a correct evaluation calibrated hardware is required, as is the use in film post-production houses. We will be producing images of rather low luminance, so they will look very different on a computer screen depending on the 'brightness' setting of the monitor, on how well lit the room is, etc. As an approximation only, we recommend the reader to set the brightness level to 50%, the monitor being in a well lit room with no direct light on the screen.

Let us comment on each of these steps now. For a very thorough coverage of Color Science, and an in-depth exposition of the concepts mentioned in this article, we refer the interested reader to the excellent treatise (Wyszecki and Stiles, 1982).

Step 1. Estimation of reflectance values. Let us say we have an object with reflectance $\beta(\lambda)$, where λ is the wavelength, and it is lit with an illuminant with spectral power $S(\lambda)$. In the XYZ color model of the CIE (Wyszecki and Stiles, 1982), the tristimulus values X, Y and Z of an object-color stimulus are obtained with these equations:

$$\begin{aligned} X &= k \int \beta(\lambda) S(\lambda) \bar{x}(\lambda) d\lambda \\ Y &= k \int \beta(\lambda) S(\lambda) \bar{y}(\lambda) d\lambda \\ Z &= k \int \beta(\lambda) S(\lambda) \bar{z}(\lambda) d\lambda \end{aligned} \tag{1}$$

Functions $\bar{x}(\lambda), \bar{y}(\lambda)$ and $\bar{z}(\lambda)$ denote the color-matching functions for a standard observer. These are experimental values tabulated by the CIE. The values X, Y and Z at each given pixel are known, we get them by converting the original R, G and B values.

The values for $S(\lambda)$ are not known precisely, unless they were measured when the original image was taken. In our model we have assumed for $S(\lambda)$ the properties of the

CIE standard illuminant D_{65}, which corresponds to a phase of natural daylight. The only unknowns are then the reflectance values. Since we want to substitute the daylight illuminant $S(\lambda)$ with a nighttime illuminant $S'(\lambda)$, we must first compute $\beta(\lambda)$. If we take just three wavelengths λ_1, λ_2 and λ_3, we can obtain a very rough estimate of $\beta(\lambda_1)$, $\beta(\lambda_2)$ and $\beta(\lambda_3)$ by solving the 3x3 system of equations we get from (1) if we replace the integrals with discrete summations of three terms. This is a very crude and simple approximation, that nonetheless produces good enough results: in figure 3 we have tested two different CIE standard illuminants without producing great perceptual changes in the simulated night scene. For this step we could use much better reflectance models like those in (Maloney, 1986) or (Slater and Healey, 1997).

Once we know the reflectance values, we substitute $S(\lambda_i)$ with $S'(\lambda_i)$, $i = 1,2,3$. We have taken λ_i near the values for monochromatic red, green and blue. In our experiments we have used for $S'(\lambda)$ the experimental values obtained in (Massey and Foltz, 2000). Surprisingly, if we only change the illuminant, the image colors get warmer. This is due to the fact that the night light $S'(\lambda)$ has more power in the long wavelengths and less in the short ones. Therefore, the perceived shift towards blue of colors under natural night light conditions is due only to the human night vision system.

Step 2. Modification of chromaticity. The perceived chromaticity depends greatly on the illumination level. Also, as we decrease the illumination level, the colors become less saturated. The experimental data in (Hunt, 1952), (Wyszecki and Stiles, 1982) and (Stabell and Stabell, 2002) show how monochromatic lights of different wavelengths are seen with evolving chromaticities as the surrounding luminances change. We use these data to modify accordingly the color matching functions $\bar{x}(\lambda)$, $\bar{y}(\lambda)$ and $\bar{z}(\lambda)$. This is nothing novel, see for instance the excellent works (Ferwerda et al., 1996), (Durand and Dorsey, 2000), (Ward et al., 1997).

Step 3. Modification of luminance. The spectral luminous efficiency function $V'(\lambda)$ measures how brightness is perceived as a function of wavelength. It depends on the illumination level, and we get it from tabulated experimental data in (Wyszecki and Stiles, 1982). From $V'(\lambda)$ the luminance L is computed.

There are two values to be set concerning luminance at night, its mean value and variance. The mean value is set by the specification of desired ambient luminance by the user.

The choice of variance sets the maximum brightness, which is crucial if we have artificial light sources in the scene, as we will see in section 3. While the variance is set by default, the user may modify it.

We proceed in this manner and compute:

$$L'_0 = \int \beta(\lambda) S'(\lambda) V'(\lambda) d\lambda$$

Then we impose the selected mean and variance for our modified luminance L':

$$L' = \frac{L'_0 - \mu}{a} \frac{b}{\mu} + b$$

where b will be the desired mean at night, μ is the mean of L'_0 (the quotient $\frac{b}{\mu}$ allows the change of units), and a controls the variance ratio of L'_0 and L' (i.e. it allows us to set any given variance for L'). Unless otherwise stated, we use a value of $a = 1$. The user may modify a to increase the variance and make artificial lights brighter, for instance.

Step 4. Modification of contrast. Because the eye adjusts to its surroundings, human sensitivity to contrast depends on the adaptation luminance (Tumblin and Rushmeier, 1993). Contrast in our night image then must be different than in the original daylight scene. We can achieve this in two ways, either by approximating the eye's performance or by simulating the use of a given type of photographic film.

To approximate the eye's performance several models have been proposed (Tumblin and Rushmeier, 1993; Ward et al., 1997; Ferwerda et al., 1996). We have implemented in our experiments a modification of the tone reproduction operator of Ward et al. (Ward et al., 1997), where they combine the rod and cone sensitivity functions and build a five-interval piecewise approximation for $\Delta L_t(L_a)$.

In their notation, ΔL_t is the 'just noticeable difference' at the given adaptation level L_a. We compute the resulting luminance L_n from the real-world luminance L_{rw} with $L_n = \frac{\Delta L_t(L'_a)}{\Delta L_t(L_{arw})} L_{rw}$, where L'_a and L_{arw} are actually 8-neighbor local averages of L' and L_{rw} respectively.

If we choose to simulate the use of a given type of photographic film with a characteristic curve of γ_n, then our night luminance L_n will be approximated as $L_n = cL'(L^{\frac{1}{\gamma}})^{\gamma_n}$, where L' stands for the original L after modification in the previous steps. If γ is not known, we just assume $\gamma = 1$ and choose γ_n to achieve visually pleasing results. (These equations may be refined at will with more accurate descriptions of the characteristic curves of the film.)

Step 5. Loss of acuity: Diffusion. Visual acuity is the ability of the eye to see fine detail. The highest level of acuity is achieved at photopic levels and it decreases as the background luminance diminishes. But it also depends on contrast: increasing the level of contrast increases the resolution at a given luminance level. In previous work (Ferwerda et al., 1996) this is modeled as isotropic diffusion. A 2D Gaussian filter is applied to the image, the radius of the Gaussian depending on the adaptation luminance. The idea is to cancel high spatial frequencies, following experimental data relating maximum visible spatial frequency (of a high contrast grating) and adaptation luminance. However, as local luminance is not taken into account, the resulting images seem unrealistic since they evoke the effects produced by an out of focus camera. In (Durand and Dorsey, 2000) there is a small correction to the computation of the Gaussian's size, but the procedure is basically the same. In (Ward et al., 1997) the authors approximate convolution with a Gaussian of explicitly varying radius: we shall prove that this approach causes ringing to appear in the resulting image. In (Thompson et al., 2002) the authors choose a spatial filtering approach, performing low-pass filtering followed by sharpening. The results are definitely better than with Gaussian blurring, but artifacts like ringing are present, and there are several parameters that are set subjectively.

Our contribution is the following: we will introduce a novel diffusion PDE that models the spatial summation principle, takes luminance into account and respects contrast, produces no ringing, is stable, very easy to implement and fast. We shall start by discussing a set of axioms modeling the loss of visual acuity, which will lead to a family of nonlinear diffusion equations which take into account local luminance and respect contrast.

Psychophysical and physiological experiments cited in (Cornsweet and Yellott, 1985) show that neighbouring photoreceptors in the retina interact accordingly to the level of illuminance at each point. That is, the light perceived at a single point in the retina not only creates an excitation at the photoreceptor at this site, it produces a lateral excitation as well and all of them are combined additively. This process is called spatial summation and the extent of the area of summation varies inversely with the local illuminance. Let I be the input image and let us denote by $O(I)$ the output image. The following set of axioms is inspired from (Cornsweet and Yellott, 1985).

(i) Each input point (x,y) contributes with a non-negative point spread value to every output point (p,q), the size of this contribution depending on the intensity value $I(x,y)$ and the distance from (x,y) to (p,q). Thus the point spread function has the form $S((x,y),(p,q),I)$ and this gives the contribution from (x,y) to (p,q) when the input intensity at (x,y) is I.

(ii) S is nonnegative, spatially homogeneous and circularly symmetric, hence $S = S(d^2,I)$, where $d^2 = (x-p)^2 + (y-q)^2$.

(iii) The effective area covered by the PSF around each input point varies inversely with the intensity at that point, which can be translated into the relation

$$S(d^2,I) = Q(I)S(Q(I)d^2,1)$$

where $Q(I)$ is an increasing function of I.

(iv) If we write $S(r^2) = S(r^2,1)$ then we normalize the integral of $S(x^2+y^2)$ over the (x,y) plane to be equal to 1, i.e.,

$$\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} S(x^2+y^2)\,dx\,dy = 1.$$

The output is given by

$$O(I)(p,q) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} I(x,y)S(d^2,I(x,y))\,dx\,dy. \tag{2}$$

In (Cornsweet and Yellott, 1985), the output was interpreted as the response of a retinal cell and it was taken as $O(I)(p,q) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} S(d^2,I(x,y))\,dx\,dy$. Notice its difference with (2) in which the output of the filter is an intensity image given by a convolution with a PSF with intensity depending variance. As in (Cornsweet and Yellott, 1985), our main example of point spread function will be the Gaussian, but we could also choose a radial function with zero first-order moments and finite second order moments. To relate $Q(I)$ with the variance of the Gaussian function, we use the notation $Q(I) = \frac{1}{\sigma^2(I)}$. Then we may write

$$S(d^2,I(x,y)) = \frac{1}{2\pi\sigma^2(I)} exp^{-\frac{[(x-p)^2+(y-q)^2]}{2\sigma^2(I)}} \tag{3}$$

Arguing as in (Cornsweet and Yellott, 1985), we observe that the spatial summation (2) does *not* satisfy the maximum principle. Indeed, if $I(x,y)$ is an step function with values

I_0 when $x < 0$ and $I_0 + D$ when $x \geq 0$, then we may find a positive value Θ such that $O(I)(p,q) > I_0 + D$ if $p > \Theta$. In other words, ringing is *guaranteed* to appear if we use directly this spatial summation, or if, equivalently, we perform convolution with Gaussians of varying radius, as in (Ward et al., 1997).

We can try to restore the maximum principle (and its locality) if we know the infinitesimal action of the filter. For that following (Alvarez et al., 1993), we introduce an scale parameter $t > 0$, and write $t\sigma(I)$ instead of $\sigma(I)$. Then the asymptotic expansion of $O(I,t)(p,q)$ around $t = 0$ is given by the following result.

THEOREM 1. *We have* $O(I,t)(p,q) = I(p,q) + Ct^2\Delta(I\sigma^2(I))(p,q) + o(t^2)$.

If we write the above expansion in the form $\frac{O(I,t)(p,q)-I(p,q)}{Ct^2} = \Delta(I\sigma^2(I))$ and let $t \to 0^+$, we obtain the following nonlinear diffusion equation to model the loss of acuity

$$I_t = \Delta(I\sigma^2(I)). \tag{4}$$

From the mathematical point of view, existence and uniqueness results for initial data $I_0 \in L^1(\Omega) \cap L^\infty(\Omega)$, together with a comparison principle for (4) have been proved in (Bénilan and Crandall, 1981) when the function $\varphi(r) = r\sigma^2(r)$ is continuous and increasing in $\Re$. We stress the fact that going from (2) to (4) permits us to guarantee the maximum principle and, therefore, no ringing behavior is exhibited by solutions of these equations. The usual rigorous derivation of this would be based on the computation of the iterates $O(I, \frac{t}{n})^n$ and its limit as $n \to \infty$, though the usual convergence Theorem (Guichard and Morel, 2000) does not apply and the convergence of the iterates $O(I, \frac{t}{n})^n$ is an open question.

To write some particular instances of (4) we write it as $I_t = \mathrm{div}(\varphi'(I)\nabla I)$. Then we choose $\varphi'(I) = \frac{1}{(1+\alpha I)^\beta}$ with $\alpha, \beta > 0$ as generic models for φ which exhibit a power behaviour for large intensities and are differentiable at $I = 0$. Thus, we shall consider

$$I_t = \nabla \cdot \left(\frac{\nabla I}{(1+\alpha I)^\beta} \right) \tag{5}$$

The parameter α controls the level of diffusion. In equation (5), the anisotropy of the diffusion is controlled by the local luminance values. Pixels with high luminance values are diffused less than pixels with low luminance. Recall that usually the anisotropy is controlled by the magnitude of the *gradient* (see (Tumblin and Turk, 1999) and the seminal work by Perona and Malik (Perona and Malik, 1990), a model in which images are smoothed while preserving edges). The level of diffusion is also controlled by parameter β, the diffusion is reduced as we increase the value of β, this effect being more pronounced for high luminances. These facts have been checked experimentally. In our algorithm we set $\beta = 1$. Figure 1(a) shows the effect of applying the PDE (5) to square grating signals of same contrast but different mean luminance (contrast is taken as the quotient $(I_{max} - I_{min})/I_{mean}$). The result proves that the square grating of higher luminance is less diffused than the lower one. On the other hand, figure 1(b) shows the diffusion effect on two square grating functions of different contrast but with the same mean luminance. The diffusion respects contrast: the left-hand square grating has higher contrast than the right-hand one, before and after diffusion. Therefore, the square grating on the left is perceived as sharper. In conclusion, we have experimentally checked that the behavior exhibited by solutions of equation (5) is consistent with experimental data on human vision acuity showing that spatial summation is inversely proportional to luminance and respects contrast. Furthermore, the results are free from ringing and other visual artifacts that other approaches bring. We compare in figure 2 our approach to different techniques to simulate the loss of acuity at night. Notice that both the convolution with gaussians of varying radius (Ward et al., 1997) and low pass filtering with sharpening (Thompson et al., 2002) produce different types of artifacts. If we convolve with a Gaussian of fixed radius depending on the adaption luminance (Durand and Dorsey, 2000) artifacts do not appear but the result has the effect of an out of focus blur.

For simplicity we apply equation (5) to each of the three color components separately, though an equivalent vector-diffusion equation could be devised. The discretization in time is performed with an explicit forward Euler scheme whereas an alternating forward-backward scheme is used for the space discretization.

We have obtained experimentally the time of diffusion T necessary to lose at each level of darkness the details whose frequency is above the Highest Resolvable Spatial Frequency in accordance with data from Shaler in (Ferwerda et al., 1996). Firstly, we have chosen three different luminances and we have created three images of a square wave grating with image dimensions corresponding to the width subtended by one degree of arc at a viewing distance of one meter. Each of these images contains a number of cycles just above the maximum number of cycles detectable at that level of luminance. The maximum value of the image is fixed to 255 for a log luminance of 3 and we decrease it proportionally to the decrease in log luminance. Then, we fix $\alpha = 0$ (isotropic diffusion) and $\Delta t = 0.1$ (below 0.25 which is the CFL stability condition for the Perona Malik equation (Perona and Malik, 1990)) and we find the necessary number of iterations to achieve a uniform image at a distance of one meter. Finally, we interpolate linearly the number of iterations between these points and we obtain the following expression for the number of steps:

$$steps = \begin{cases} 12 - 36log(L) & \text{if } log(L) < 0 \\ 12 - 6.4log(L) & \text{if } 0 \leq log(L) < 1.875 \\ 0 & \text{if } log \geq 1.875 \end{cases}$$

We have set α to 0.01, obtaining very good results for natural images in a wide range of ambient luminances. Please note that this value of α is fixed in the algorithm and thus it is not a parameter that the user has to change. The robustness of the equation makes it suitable for video sequences, no temporal artifacts appear (see examples in http://www.tecn.upf.es/˜mbertalmio/day4nite).

Figure 4 shows how fine details are lost as the luminance level decreases. Also, for any given image and luminance level, more detail is lost in darker regions than in light ones. Notice how the achieved effect of loss of acuity is very different from an out-of-focus blur. In particular, in figure 4, as the luminance decreases it becomes harder and harder to read the numbers on the wall, or the text in banners, books and cardboard boxes, just as it happens to our eyes when the light grows dim. But *pronounced* edges are preserved, as we can see in the dark bands on the wall, or the white sheets of paper hanging from the tables in this same figure.

3. Examples

Figures 3 to 6 show several results, for different images, night illumination levels and contrast-modification methods. Notice how these images look realistic. In particular, notice how colors have become less saturated but we may still tell them apart, they are not predominantly different shades of blue as we would get with conventional Day for Night. Brightness and contrast are what we would expect in a night scene, objects do not have an unreal illumination. Realism is enhanced by the controlled loss of resolution, which blurs small (and not too bright) details, as our eyes do at night. We show in figure 3 that the algorithm is robust to the natural illuminant that we assume: the results do not vary much using two different CIE standard illuminants D_{65} and D_{75}, which correspond to natural daylight at different color temperatures.

Our algorithm has been developed with the assumption that all light in the scene is natural, i.e. that the illuminant is one for the whole image. We are currently working on how to circumvent this constraint, so we can introduce artificial light sources in our images. The problem is that it is very hard to approximate, at each pixel location, the interaction between different light sources, with different intensities and spectral power. Figure 6 shows a test for one image of this sort, where we have assumed that the highest luminances in the scene correspond to the artificial light source. We have increased the luminance there modifying the original variance by a factor greater than 1. The results are more realistic but this method can fail at points with high luminance but which do not correspond to light sources (see white line on the road in figure 6).

If the original image presents a cloudy sky, as in figure 4, we cannot achieve a dark sky in the night scene. If this is a problem we could choose to avoid showing the sky when shooting, as done in traditional Day for Night. We could also explore a way to segment the sky and treat differently the pixels in that region.

The whole process takes less than ten seconds in a regular PC for a 512x768 24 bits RGB image. Within that execution time, seven seconds are dedicated to the diffusion process in the three color components in an example where seven iterations of the equation are required. The use of a GPU (Graphics Processor Unit) may speed-up the anisotropic diffusion process by an order of magnitude (Bertalmío et al., 2004), bringing the algorithm closer to real-time. We have also verified that the parallelization of the code does not create artifacts. By constructing a color LUT (Look-Up Table) the speed may be increased greatly from our current implementation, where we deal with each pixel separately. Also in moving pictures there is great space and time redundancy, another source for speed-ups.

4. Conclusion and future research

We have introduced a digital Day for Night algorithm that achieves realistic results. Our algorithm performs modification of the spectrum for the night illuminant, desaturation of the colors, brightness modification according to wavelength, contrast modification according to luminance adaptation levels, and non-uniform and non-linear loss of resolution. We use a set of simple equations, based on real physical data and visual perception experimental data. To simulate the loss of resolution we have introduced a novel diffusion equation, which is well-posed, has existence and uniqueness results, and is also monotonicity preserving, so no ringing may occur. The robustness of the equation makes it suitable for video sequences, no temporal artifacts appear. The user only has to provide the original day image and the desired level of darkness of the result. The whole process from original day image to final night image takes a few seconds, all the computations being local, but optimizations could easily speed up the process in an order of magnitude.

The main limitation of our algorithm is that it has been developed with the assumption that all light in the scene is natural, i.e. that the illuminant is one for the whole image. We are currently working on how to circumvent this constraint. The input images are in RGB format coming from digital video or obtained from a scanned film. Part of future research is to include emulations of the film developing process, and to reformulate our algorithm in terms and units that cinematographers use.

5. Acknowledgements

The authors especially wish to thank Pere Freixa for all his cooperation, support and helpful comments. Many thanks go also to Joan Abadía, Rocío García, Ricardo Gadea, Hugo Peluffo and Josep Blat for their help. The first author acknowledges partial support by grant number FP2000-5801. The second author acknowledges partial support by Programa Ramón y Cajal. The third author acknowledges partial support by the DURSI (Generalitat de Catalunya). We all acknowledge partial support by PNPGC project, reference BFM2003-02125 and by the European Commission Project IST-2001-34337.

This work is dedicated to Andrés Solé.

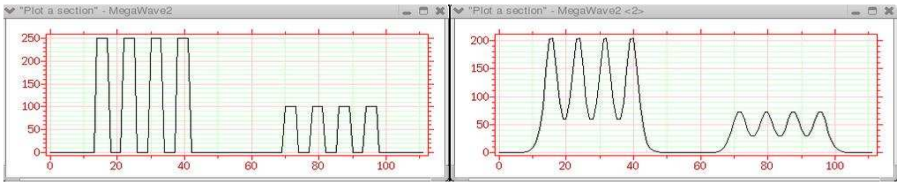

(a) Diffusion on two square grating functions of same contrast.

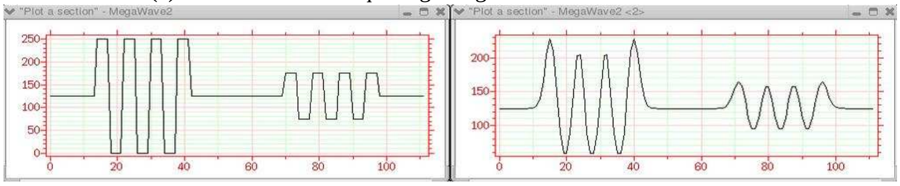

(b) Diffusion on two square grating functions of different contrast.

Figure 1. Loss of acuity simulation via PDE's: diffusion depends on luminance. Observe that, after diffusion, the left grating is still more contrasted than the right one (25 iterations, dt=0.1, α=0.01).

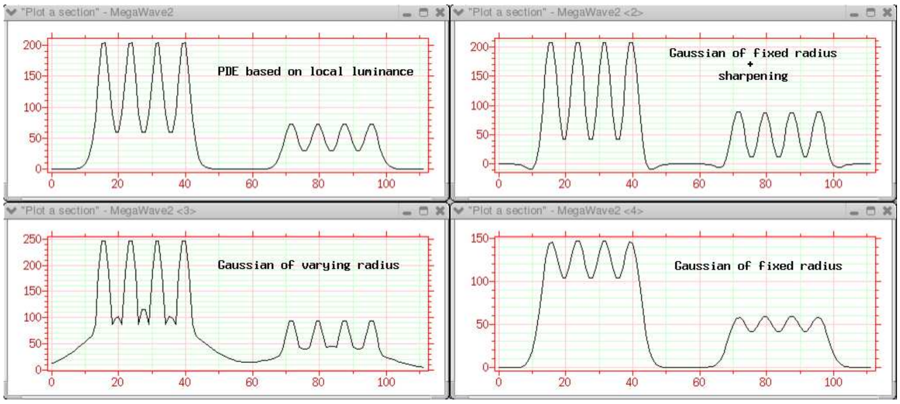

Figure 2. Comparison of different diffusion methods simulating loss of acuity at night. Notice that our method (top left) produces no artifacts, neither the Gaussian of fixed radius (bottom right) but this approach does not simulate the spatial summation principle depending on local luminance.

References

Alvarez, L., F. Guichard, P. L. Lions, and J. M. Morel: 1993, 'Axioms and fundamental equations of image processing'. *Archive Rat. Mech. and Anal.* pp. 200–257.

Bénilan, P. and M. Crandall: 1981, 'The continuous dependence on φ of solutions of $u_t - \Delta\varphi(u) = 0$'. *Indiana Math. J.* **30**(2), 161–177.

Figure 3. Left, original scene. Next, our results for two night simulated scenes at 1 log cd/m^2 with different assumptions for the day illuminant: D_{65} (middle) and D_{75} (right).

Figure 4. Some results with decreasing values of ambient luminance: 1, 0.6, 0.3, 0.1 and -0.1 log cd/m^2 , 5, 8, 10, 11 and 15 iterations of diffusion respectively from left to right and from top to bottom.

Bertalmío, M., P. Fort, and D. Sánchez-Crespo: 2004, 'Real-time, accurate Depth of Field using Anisotropic Diffusion and Programmable Graphics Cards'. In: *Proceedings of 2nd 3DPVT*.

Cornsweet, T. N. and J. I. Yellott: 1985, 'Intensity-independent spatial summation'. *Journal of Optical Society of America* **2**(10), 1769–1786.

Durand, F. and J. Dorsey: 2000, 'Interactive Tone Mapping'. In: *Proc. Eurographics Workshop on Rendering*.

Ferwerda, J. A., S. N. Pattanaik, P. Shirley, and D. P. Greenberg: 1996, 'A Model of Visual Adaptation for Realistic Image Synthesis'. In: *Proceedings of SIGGRAPH 1996*. pp. 249–258.

Guichard, F. and J. Morel: 2000, *Image Iterative Smoothing and P.D.E.'s'*. Book in preparation.

Hunt: 1952, 'Light and Dark Adaptation in the Perception of Color'. *J. Optical Soc. America A* **42**(3), 190.

Lumet, S.: 1995, *Making Movies*. Alfred A. Knopf.

Maloney, L. T.: 1986, 'Evaluation of linear models of surface spectral reflectance with small numbers of parameters'. *J. Optical Soc. America-A* **3**(10), 1673–1683.

Figure 5. Original scene (left), emulating night vision (middle), emulating a given film stock (right).

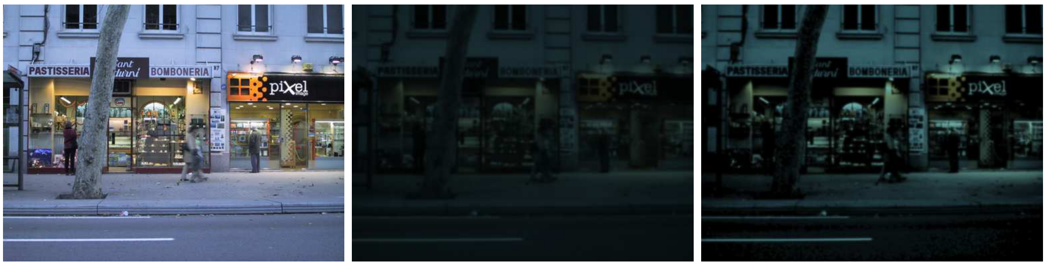

Figure 6. Original scene (left), result with variance ratio $a = 1$ (middle), increasing variance of luminance with $a = 0.1$ (right).

Massey, P. and C. B. Foltz: 2000, 'The Spectrum of the Night Sky Over Mount Hopkins and Kitt Peak: Changes After a Decade'. *Publications of the Astronomical Society of the Pacific* **112**(566).

Perona, P. and J. Malik: 1990, 'Scale-space and edge detection using anisotropic diffusion'. *IEEE TPAMI* **12**(7), 629–639.

Slater, D. and G. Healey: 1997, 'Using a spectral reflectance model for the illumination-invariant recognition of local image structure'. *IEEE TPAMI* **19**(10).

Stabell, B. and U. Stabell: 2002, 'Effects of rod activity on color perception with light adaptation'. *Journal of Optical Society of America* **19**(7), 1249–1258.

Thompson, W., P. Shirley, and J. Ferwerda: 2002, 'A Spatial Post-Processing Algorithm for Images of Night Scenes'. *Journal of Graphics Tools* **7**(1), 1–12.

Tumblin, J. and H. Rushmeier: 1993, 'Tone Reproduction for Realistic 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: *Siggraph 1999, Computer Graphics Proceedings*. pp. 83–90.

Ward, G., H. Rushmeier, and C. Piatko: 1997, 'A Visibility Matching Tone Reproduction Operator for High Dynamic Range Scenes'. *IEEE Trans. on Visualization and Computer Graphics* **3**(4).

Wyszecki, G. and W. S. Stiles: 1982, *Color Science: Concepts and Methods, Quantitative Data and Formulae (2nd ed.)*. New York: John Wiley & Sons, Inc.