

Visual Acuity in Day for Night

Gloria Haro, Marcelo Bertalmío and Vicent Caselles
Departament de Tecnologia
Universitat Pompeu Fabra
Pg. de Circumval.lació, 8; 08003 Barcelona, Spain
e-mail: { gloria.haro,marcelo.bertalmio,vicent.caselles } @upf.edu

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. These problems stem from the fact that this Day for Night process is exclusively optical (filters, film exposure), and hence it can't provide the modifications that account for how we see the world at night: different spectrum for the night illuminant, desaturation of colors, brightness modification according to wavelength, contrast modification according to luminance adaptation levels, and non-uniform and non-linear loss of resolution. 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 both luminance and contrast into account, 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.

1 Introduction

In his book 'Making movies' [10], the great American director Sidney Lumet points out the difficulties of night shooting on location. Apart from the inherent lack of control that locations have (as opposed to studios), night shooting requires that *everything* be lit artificially. This may be a big problem if the location covers a wide area. In some lo-

cations the terrain may make it difficult to bring the truckloads of lamps and generators that the shooting requires. Because they make a lot of noise, generators are placed far from the set so as not to interfere with the sound department. Very long cables have to be laid from the lights to the generators, and the rigging crew and electricians have much more extra work. If the night shooting continues for weeks, cast and crew get exhausted. And on the days when night shooting takes place, there is time for very little or no day shooting at all, which may complicate the schedule. 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. They are dimmer than with daylight, but somehow we still can see everything that is in the scene. Moreover, no detail is lost: if there is a sign with small fonts, we can read it as if it were noon.

The problem is that with mere optical means (blue filters, under-exposure) we can not expect to reproduce all the modifications that account for how we see the world at night. Firstly, let us assume that no artificial light source is present in our scene, i.e., that all the light is supposed to be coming from the moon, the sky and the stars. While it is true that at night we perceive blue objects brighter, this is *not* due to the light at night being bluer. Actually, light at night has a power spectrum with a stronger red component than daylight, and with less blue [9]. It is the human visual system that 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. Currently our implementation only deals with natural light sources, but we hint at a way to incorporate artificial lighting to our scenes. 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. For the simulation of the loss of visual acuity, we introduce a novel diffusion PDE that models the spatial summation principle [4], takes both luminance and contrast into account, produces no ringing, is stable, very easy to implement and fast. We start by replacing daylight illumination with night-like illumination, a procedure which to the best of our knowledge is novel when addressing this problem.

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. While the user would typically only care about the final image, we present here all the steps separately. This may help the user to modify some steps to obtain a more expressive result. The aim is to achieve realistic and visually pleasing results, even if they are not entirely compliant with the models of human visual perception.

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. For this we assume that the day scene has been lit with daylight and we use the standard illuminant D_{65} to approximate it. Then we replace the D_{65} with an estimation of the spectrum of the night sky. This is, to the best of our knowledge, a novel procedure in the context of this problem. 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, since the luminous efficiency function depends on the illumination level of the scene. 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 [16, 19, 7]. 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 [4].

One note concerning the evaluation of the results. We will be producing images of rather low luminance, so they will look very differently on a computer screen in a well lit room than on a dimly lit room. Another crucial factor

for the evaluation is the 'brightness' setting of the monitor. 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. This would be an approximation only, for correct evaluation calibrated hardware is required.

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 by Wyszecki and Stiles [20].

2.1 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 [20], the tristimulus values X, Y and Z of an object-color stimulus are obtained with these equations:

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

where the factor k is defined so that the Y tristimulus for the *perfect reflection diffuser* ($\beta(\lambda) = 1.0$ at all wavelengths) is equal to 100. 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 (i.e. we go from the RGB to the XYZ color model).

Care must be taken in that the RGB image has been acquired non-linearly from real-world luminances [16]. Each type of photographic film has a characteristic curve that determines how luminances in the real world (L_{rw}) are encoded as transparencies (T) in the film: $T = a_1 L_{rw}^\gamma$, where the constant factor a_1 depends on the film speed and the choice of exposure time, lens and camera aperture. This function $T(L_{rw})$ is valid within a certain interval of luminances, outside which the value T is 'clipped'. Since we need to work with real-world luminances, we must undo the transformation: $T' = a_2 T^{\frac{1}{\gamma}}$. For this we need the contrast sensitivity specification of the film stock that was used to shoot the original image. If we don't have this data, we can nonetheless obtain good results by assuming a value for γ and selecting a value for the 'night γ ' that gives a visually convincing result, as is done in subsection 2.4.

Returning to equation (1), the values for $S(\lambda)$ are not known precisely, unless they were measured when the original image was taken. But it is safe to assume 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 sys-

tem of equations we get from (1) if we replace the integrals with discrete summations of three terms. Once we know these 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, so we may also call them λ_r , λ_g and λ_b .

From now on in our work, we will be dealing with this discrete set of three values per continuous function involved. This is a very crude approximation, motivated by the fact that we only start with three values per pixel, the R, G and B values. In our experiments we have used for $S'(\lambda)$ the experimental values obtained in [11]. If we only change the illuminant, the image colors get warmer. This, as we mentioned above, is due to the fact that the night light $S'(\lambda)$ has more power in the long wavelengths and less in the short ones (i.e. more 'red' and less 'blue').

2.2 Modification of chromaticity

The perceived chromaticity depends greatly on the illumination level. As we decrease the illumination level, the colors become less saturated. A color that is very vivid under daylight seems less and less vivid as the illumination decreases. This property is very difficult to emulate directly on film, with techniques ranging from the pre-exposing of film [10] to the 'bleach by-pass' at the developing stage [6]. The experimental data in [8], [20] and [13] 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 [7], [5], [19].

2.3 Modification of luminance

There are two kind of photoreceptors in the retina, the rods and the cones. Under low light (scotopic) conditions only the rods are active. Since there is only one type of rods, nocturnal vision is monochromatic. At daylight (photopic) light levels the rods become saturated and all the visual information comes from the cones. There are three types of cones, each one is tuned to a certain wavelength range. *Trichromatic generalization* explains how color stimuli can be expressed as additive mixtures of three fixed primary stimuli. At intermediate illumination levels (mesopic conditions) both rods and cones are active. The spectral luminous efficiency functions $V(\lambda)$ and $V'(\lambda)$ are tabulated by the CIE and measure how brightness is perceived as a function of wavelength in the photopic and scotopic ranges respectively. In the mesopic range, the spectral luminous efficiency function depends on the illumination level, and we get it from tabulated experimental data in [20]. It can not be approximated by a linear combination of $V(\lambda)$ and $V'(\lambda)$. From the spectral luminous efficiency function, 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. While the variance is set by default, the user may modify it. 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.

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.

2.4 Modification of contrast

Because the eye adjusts to its surroundings, human sensitivity to contrast depends on the adaptation luminance [16]. 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 [16, 19, 7]. We have implemented in our experiments a modification of the tone reproduction operator of Ward et al. [19], 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 $L_n = \frac{\Delta L_t(L_a)}{\Delta L_t(L_{a_{rw}})} L_{rw}$, where L'_a and $L_{a_{rw}}$ 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 $L_n = cL'(L'^{\frac{1}{\gamma_n}})^{\gamma_n}$, 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.)

An example, in figure 2, compares both methods (in this figure we have also simulated the loss of acuity, see section 2.5.) Figure 2 (right) shows the result of simulating a photographic film. To most observers, the result may be visually more pleasing than that obtained with an emulation of the eye's performance (figure 2 (middle)). The reason for

this is that we are accustomed to a certain “look” of night images on photographs and movies, which we usually prefer over “real” looking night images, that may appear as too dark or without enough contrast. On the other hand, if the original day image has a sharp contrast, the “night” result may be very convincing. In any case, the user may select in our algorithm the method that suits him/her best.

2.5 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 [7] 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 neither contrast nor local luminance are taken into account, the resulting images seem unrealistic since they evoke the effects produced by an out of focus camera. In [5] there is a small correction to the computation of the Gaussian’s size, but the procedure is basically the same. In [19] 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 [15] 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 will be the following. We will prove that there is a family of PDE’s that diffuse a given image adequately matching the loss of acuity in night vision. The diffusion process takes both local luminance and contrast into account. And unlike previous methods, ringing is guaranteed not to occur. When working on a sequence of images, no temporal artifacts appear either. The implementation of these PDE’s is straightforward, and they run fast since the number of iterations is very low. These PDE’s are called Fast Diffusion Equations, and are related to Porous Medium Equations. The method we introduce here could also be used to simulate loss of acuity on images shot at night, since film and cameras can not emulate this human vision process.

Psychophysical and physiological experiments cited in [4] 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 in-

versely with the local illuminance.

Indeed, in [4] the authors start from a set of axioms for intensity dependent spatial summation to determine the point spread function relating the input image I to the output image $O(I)$. The basic idea is that is that 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 . Then they assumed

- (i) S is nonnegative,
- (ii) S is 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 . Indeed, in [4], the authors took $Q(I) = I$.

- (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.,

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

Then the main example used by the authors is the Gaussian function. 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

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

and

$$O(I)(p,q) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} I(x,y)S(d^2, I(x,y)) dx dy. \quad (3)$$

Then the authors discussed the main properties of this filter [4]. Apart from its expected effect of improving spatial resolution as illuminance increases, this mechanism also yields edge-enhancement (Mach bands) and other band-pass filter effects. In fact, working with a step function as input I , and operating on the integral $O(I)$ we obtain the following result:

Theorem 1 *We have $O(I)(p,q)$ as in eq.(3), and $I(x,y)$ is a step function valued I for $x < 0$ and $I+D$ for $x \geq 0$. Then there exists a positive real number Θ s.t. $O(I)(p,q) > I+D$ if $p > \Theta$.*

In other words, ringing is *guaranteed* to appear if we use this approach.

We will propose now a PDE that performs this spatial summation process. The main assumption in the set of axioms above is that the effective radius of the PSF depends inversely on the intensity I . Hence, we shall assume that

$$\sigma(I) \text{ is a decreasing function of } I.$$

To study the infinitesimal action and the local effect of the filter, we shall proceed as usual [1], we introduce an scale parameter $t > 0$, and write $t\sigma(I)$ instead of $\sigma(I)$. Then we compute the asymptotic expansion of $O(I,t)(p,q)$ around $t = 0$ to obtain the following result.

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

Thus the underlying PDE is the nonlinear diffusion equation

$$I_t = \Delta(I\sigma^2(I)) \quad (4)$$

As examples we may take $\sigma^2(I) = I^{-\beta}$, $0 < \beta < 1$, or $\sigma^2(I) = \frac{\log(1+\alpha I)}{I}$, $\alpha > 0$. If the function $\varphi(r) = r\sigma^2(r)$ is continuous and increasing in \mathfrak{R} , equation (4) with Neumann boundary conditions is well posed for any initial condition $I_0 \in L^1(\Omega) \cap L^\infty(\Omega)$, and satisfies a comparison principle. Therefore there is no ringing behavior for these equations. For the mathematical treatment of this kind of equations we refer to [2], [3, 18] and references therein.

As a particular case of (4), we take $\sigma^2(I) = \frac{\log(1+\alpha I)}{I}$, $\alpha > 0$, which gives:

$$I_t = \nabla \cdot \left(\frac{\nabla I}{1+\alpha I} \right) = \Delta(\log(1+\alpha I)) \quad (5)$$

where I is a function that represents the level of luminance at each site, and α is a tunable parameter that controls the level of diffusion. In equation 4, 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 [17] and the seminal work by Perona and Malik [12], a model in which images are smoothed while preserving edges). Furthermore, experimental data on human vision acuity show that spatial summation is not only inversely proportional to luminance but also inversely dependent on contrast. This is also the behavior exhibited by solutions of equation 5, as we have experimentally checked in the case of square waves (see experiment in figure 1). We wonder if there is a connection of the above models with Weber's, Fechner's, or Stevens' Law [20].

For simplicity we apply this equation to each of the three color components separately, though an equivalent vector-diffusion equation could be devised after [14].

The numerical implementation of each scalar diffusion equation is done with a scheme based on finite differences. If we consider the representation of a color component

value at each point of the image grid as $I_{i,j} = I(i,j)$, with $1 \leq i \leq N$ and $1 \leq j \leq M$ (N is the number of lines and M the number of columns) and the finite differences at both sides to represent the spatial derivatives, we get:

$$\nabla^+ I_{i,j} = (I_{i+1,j} - I_{i,j}, I_{i,j+1} - I_{i,j})$$

$$\nabla^- I_{i,j} = (I_{i,j} - I_{i-1,j}, I_{i,j} - I_{i,j-1})$$

The numerical scheme used is:

$$I^{n+1} = I^n + \Delta t \begin{cases} \nabla^- \cdot \left(\frac{\nabla^+ I}{1+\alpha I} \right) & \text{if } n = \dot{2} \\ \nabla^+ \cdot \left(\frac{\nabla^- I}{1+\alpha I} \right) & \text{if } n = \dot{2} + 1 \end{cases} \quad (6)$$

where I^n is the image I at time n and Δt is the time step between two iterations.

We have obtained experimentally the time of diffusion T necessary to loose at each level of darkness the details whose frequency is above the Highest Resolvable Spatial Frequency in accordance with data from Shaler in [7]. 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 average 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 [12]) 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:

$$\text{steps} = \begin{cases} 12 - 36\log(L) & \text{if } \log(L) < -0.5 \\ 12 - 6.4\log(L) & \text{if } -0.5 \leq \log(L) < 1.875 \\ 0 & \text{if } \log \geq 1.875 \end{cases}$$

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.

Equation (5) is well-posed, has existence and uniqueness results, and is also monotonicity preserving, so no ringing may occur. The robustness of the equation make it suitable for video sequences, no temporal artifacts appear (see examples in <http://www.tecn.upf.es/~mbertalmio/day4nite>).

Figure 4 show 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 2 to 4 show several results, for different images, night illumination levels and contrast-modification methods. Notice how these images look quite realistic. These results must be visually convincing by themselves, the point being that the observer does not notice that they come from day images. 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.

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 3 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 3).

If the original image presents a cloudy sky, as in figures 2 and 4, we can not 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 is in the order of a few seconds in a regular PC for a 600x800 24 bits RGB image, most of the computations being local. We have not optimized our code, but since what we are doing is basically constructing a color LUT (and then diffusing the resulting image), 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 very 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 very 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, possibly connected with Weber's law, which is well-posed, has existence and uniqueness results, and is also monotonicity preserving, so no ringing may occur. The robustness of the equation make 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. More accurate results are obtained if our algorithm is provided with the characteristic curve of the photographic film used. 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. This work was partially funded by Programa Ramón y Cajal, MCyT, Spain, and by the European Commission Project IST-2001-34337. This work is dedicated to Andrés Solé.

References

- [1] L. Alvarez, F. Guichard, P. L. Lions, and J. M. Morel. Axioms and fundamental equations of image processing. *Archive Rat. Mech. and Anal.*, pages 200–257, 1993.
- [2] Ph. Bénilan and M.G. Crandall. The continuous dependence on φ of solutions of $u_t - \delta\varphi(u) = 0$. *Indiana Math. J.*, 30(2):161–177, 1981.
- [3] E. Chasseigne and J.L. Vázquez. Theory of extended solutions for fast diffusion equations in optimal classes of data. radiation from singularities. *Archive Rat. Mech. Anal.*, 164:133–187, 2002.

- [4] T. N. Cornsweet and J. I. Yellott. Intensity-independent spatial summation. *Journal of Optical Society of America*, 2(10):1769–1786, 1985.
- [5] Frédo Durand and Julie Dorsey. Interactive tone mapping. In *Proceedings of the Eurographics Workshop on Rendering*. Springer Verlag, June 2000.
- [6] Peter Etedgui. *Screencraft: Cinematography*. Focal Press, 1999.
- [7] J. A. Ferwerda, S. N. Pattanaik, P. Shirley, and D. P. Greenberg. A model of visual adaptation for realistic image synthesis. In *Proceedings of SIGGRAPH 1996*, pages 249–258. ACM Press / ACM SIGGRAPH, 1996.
- [8] Hunt. Light and dark adaptation in the perception of color. *Journal of the Optical Society of America A*, 42(3):0–0, 1952.
- [9] H.W. Jensen, F. Durand, M.M. Stark, S. Premoze, J. Dorsey, and P. Shirley. A physically-based night sky model. In *Proceedings of SIGGRAPH 2001*. ACM Press / ACM SIGGRAPH, 2001.
- [10] Sidney Lumet. *Making Movies*. Alfred A. Knopf, 1995.
- [11] Philip Massey and Craig B. Foltz. 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), 2000.
- [12] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.
- [13] B. Stabell and U. Stabell. Effects of rod activity on color perception with light adaptation. *Journal of Optical Society of America*, 19(7):1249–1258, 2002.
- [14] B. Tang, G. Sapiro, and V. Caselles. Color image enhancement via chromaticity diffusion. *IEEE Transactions On Image Processing*, 10(5):701–707, 2001.
- [15] W.B. Thompson, P. Shirley, and J.A. Ferwerda. A spatial post-processing algorithm for images of night scenes. *Journal of Graphics Tools*, 7(1):1–12, 2002.
- [16] J. Tumblin and H. Rushmeier. Tone reproduction for realistic images. *IEEE Computer Graphics and Applications*, 13(6):42–48, 1993.
- [17] Jack Tumblin and Greg Turk. LCIS: A boundary hierarchy for detail-preserving contrast reduction. In *Siggraph 1999, Computer Graphics Proceedings*, pages 83–90, 1999.
- [18] J. L. Vázquez. An introduction to the mathematical theory of the porous medium equation. *Shape Optimization and Free Boundaries (Montreal, PQ)*, pages 347–389. NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci., Kluwer Acad. Publ., Dordrecht, 1992.
- [19] G. Ward, H. Rushmeier, and C. Piatko. A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Trans. on Visualization and Computer Graphics*, 3(4), 1997.
- [20] G. Wyszecki and W. S. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulae (2nd ed.)*. John Wiley & Sons, Inc., New York, 1982.

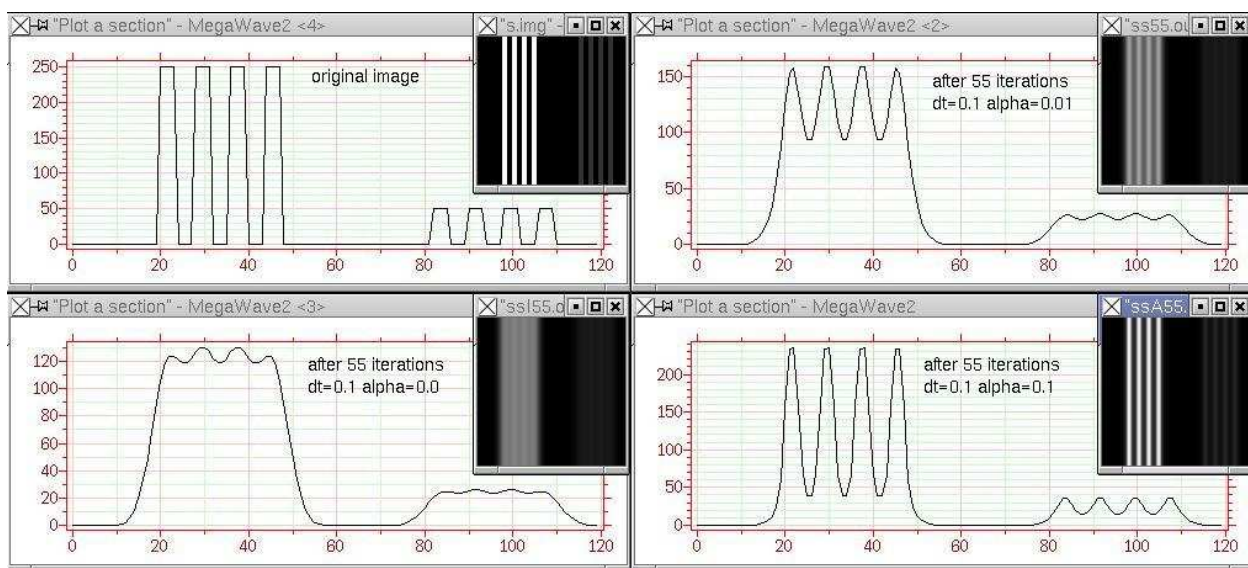


Figure 1. Loss of acuity simulation via PDE's: diffusion depends on luminance and contrast.



Figure 2. Original scene (left), emulating night vision (middle), emulating film (right).




Figure 3. Original scene (left), result with variance ratio $a = 1$ (middle), increasing variance of luminance with $a = 0.1$ (right).



Figure 4. Some night scenes 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.