Simulation of light and dark adaptation in VR

Ondrej Baranovič (ob341)

Abstract

Because of the high range of illumination we encounter in our daily lives, our visual system evolved an adaptation mechanism. This mechanism allows us to distinguish details both in direct sunlight and at night by adjusting our sensitivity to light. High dynamic range (HDR) photographs capture a wider range of luminance values than can be displayed on current consumer displays. In this project, I develop a simple model for simulating *light* and *dark adaptation* of the human visual system that can be used to display panoramic HDR photographs in a virtual reality (VR) headset. This model, based on psychophysical data, allows the user to distinguish a wider range of details by dynamically adjusting the tonemapping parameters based on the visible portion of the HDR photograph.

1 Introduction

The human visual system has evolved to support the wide range of visual conditions that are encountered in daily life. We can see both during the day when the average ambient light level approaches 10^4 cd/m^2, but also at night when the ambient light level drops to around 10^{-3} cd/m^2. However, due to the very large contrast between the two levels, our eyes cannot make out both bright and dark features at all times, and need time to adapt to the ambient light level. *Dark adaptation* is the process of adapting from a bright to a dark environment such as when entering a dark room, while *light adaptation* is the process of adapting from a dark to a bright environment such as when turning on a lamp.

Similarly, displays have a limited ability to reproduce very bright or very dark images. Current high end consumer displays can reach only up to 1000 cd/m^2 maximum brightness, and only for a small part of the screen at a time. The brightness of visual content such as photographs therefore needs to be adjusted to match the capabilities of the display. For dynamic content such as VR, the optimal brightness can vary depending on the user's input as the user looks around at different parts of the scene. For example, it is sensible to match the range of light levels in a dark room reproduced in VR to the full range of the display, but if the user then turns around to look at a bright object such as a window, its brightness would be outside the capabilities of the display. A dynamic approach could then dim the whole screen to let the user make out detail outside the window.

In this project I compare three approaches for determining the optimal brightness for displaying a panoramic HDR photograph in VR. The first is a static global approach that analyzes the whole scene and matches its brightness to the range supported by the display. The second is a dynamic approach that matches the brightness of the currently visible portion of the photograph to the display. The final approach simulates *light* and *dark adaptation* of our visual system by slowly adjusting the brightness of the displayed image over time based on the currently visible portion of the photograph.

2 Related work

In this project I examine and build upon the work of two previous papers. Ferwerda et al. [1996] proposed a computational model of visual adaptation for synthesis of realistic images. Their model is based on psychophysical experiments and, as such, it can be used to predict the appearance of features in a scene. In particular, they propose a tone reproduction operator based on threshold visibility models such that *just noticeable* differences in the scene are mapped to just noticeable differences in the image. They combined this operator with their model for *light* and *dark adaptation* to simulate the appearance of a synthetic scene at various light levels.

Wernikowski et al. [2019] continued this work by applying the tone reproduction operator to synthetic scenes displayed in VR. They proposed a colour discrimination operator to simulate our lack of colour vision at very dark illumination when only our rods are able to make out differences in a scene. They also showed that simulating realistic *dark adaptation* is not suitable for VR applications and a faster model provides a preferable experience. Neither group published their model of *light* and *dark adaptation* in sufficient detail to reproduce it.

I use Ferwerda et al.'s tone reproduction operator and Wernikowski et al.'s colour discrimination operator to display HDR panoramic photographs in VR. I first examine the results when the simulated light adaptation level is fixed for the global illumination, then adjust it based on the visible part of the scene, and finally, I propose a simple model of *light* and *dark adaptation* that works well in VR. I compare these three approaches and their suitability for VR applications, as well as the suitability of Ferwerda et al.'s and Wernikowski et al.'s operators for displaying HDR panoramic photographs.

While developing this application I found that Unity developed their own post-processing effect called "Auto Exposure" or "Adaptation" which also aims to simulate the effect of light and dark adaptation. Unlike my approach based on threshold visibility models, Unity's model is based on photographic exposure. This can be a good choice for video games, where the player can imagine that the display is provided via a simulated camera, but it does not provide a realistic experience, especially during dark-to-light adaptation. There is ev-

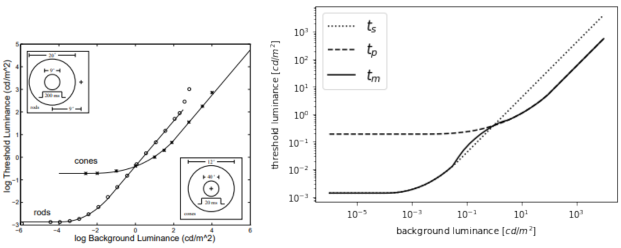

(a) Experimental data from Ferwerda et al. (b) Plot of used threshold functions

Figure 1: Threshold visibility functions.

idence that this effect was adjusted to work with VR, and I expect it would work well with abstract visual art. Unity does not appear to provide an effect to desaturates dark images in the scotopic and mesopic range.

3 Method

The VR application I produced for this project renders several HDR panoramic photographs as the skybox. It then calculates the simulated light adaptation level, applies Ferwerda et al.'s [1996] tone reproduction operator based on this state and Wernikowski et al.'s [2019] colour discrimination operator based on the scene luminance. Finally, it uses a basic tonemapping curve to transform the HDR image to the SDR colour space of the VR display.

In this section I first describe the threshold visibility function used by the tone reproduction operator, my light and dark adaptation model, and finally the tone reproduction and colour discrimination operators.

3.1 Threshold visibility

Following the two previous papers, I based my model on threshold sensitivity models. As I did not have access to the raw psychophysical data used by previous authors, I used the functions proposed by Ferwerda et al. fitting their experimental data:

$$\log t_p(L_a) = \begin{cases} -0.72 & \text{if } \log L_a \leq -2.6, \\ \log L_a - 1.255 & \text{if } \log L_a \geq 1.9, \\ (0.249\ \log L_a + 0.65)^{2.7} - 0.72 & \text{otherwise} \end{cases} \tag{1}$$

$$\log t_s(L_a) = \begin{cases} -2.86 & \text{if } \log L_a \leq -3.94, \\ \log L_a - 0.395 & \text{if } \log L_a \geq -1.44, \\ (0.405\ \log L_a + 1.6)^{2.18} - 2.86 & \text{otherwise} \end{cases} \tag{2}$$

Given the level of background illumination L_a the user is fully adapted to, the functions $t_p(L_a)$ and $t_s(L_a)$ give the threshold sensitivity for cones and rods respectively. Inspired by Ferwerda et al., I also used a parameter k to interpolate between the two functions to create a combined sensitivity function for the full range of human vision:

$$k(L_a) = \begin{cases} 0 & \text{if } L_a \leq 0.03 \\ 1 & \text{if } L_a \geq 3 \\ (0.5 \log L_a + 0.761)^{1.5} & \text{otherwise} \end{cases} \tag{3}$$

$$t_m(L_a) = t_s(L_a) \cdot (1 - k(L_a)) + t_p(L_a) \cdot k(L_a) \tag{4}$$

These functions are later used to simulate light and dark adaptation and to compute the tone reproduction parameters. A plot of Ferwerda et al.'s original data and these functions is given in fig. 1.

3.2 Adaptation

Ferwerda et al.'s tone reproduction operator is based on the background luminance adaptation of the simulated world observer L_{wa}. In this project I compared three methods for determining its value. In all three methods, I first compute the optimal adaptation luminance $\overline{L_w}$ that would be reached after waiting an infinite amount of time. I estimate the optimal adaptation luminance as the arithmetic mean of pixel luminances, or, equivalently, the total luminous intensity hitting the camera divided by the size of the display. I do not use the maximum image luminance as Ferwerda et al. because I believe the amount of light entering the eye is just as important as the brightest point.

For the first method, I precompute the arithmetic mean luminance of the whole panoramic image and use it as the optimal value. Because panoramic images are represented in an equirectangular projection, each pixel is weighted by its relative size on a sphere, $\sin \frac{\pi \cdot y}{height}$. For the second and third method, I compute the weighted arithmetic mean luminance in an area around the center of the visible portion of the image. I found that tweaking the parameters of my weight function provided advantages in different scenes, but I settled on a single weight function (plot in fig. 2) that gives acceptable results across all used images,

$$w(d) = \begin{cases} 1 - 4 \cdot d^2 & \text{if } d < 0.5 \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

where d is the distance of a given pixel from the center of the screen relative to the distance from the center to the edge of the screen:

$$d = \left\| \frac{2 \cdot x - width}{width}, \frac{2 \cdot y - height}{height} \right\| \tag{6}$$

For the first and second method, the world adaptation luminance L_{wa} is set to the average luminance $\overline{L_w}$, but for the third method I simulate light and dark adaptation. Ferwerda et al.'s chart of the time course of dark adaptation has a shape similar to the fully adapted threshold visibility data (fig. 1a). Based on this, I assumed that dark adaptation can be simulated by a *simple exponential smoothing* function, and

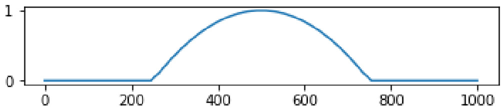

Figure 2: Weight of pixels along the center row of a 1000px wide screen when computing $\overline{L_w}$.

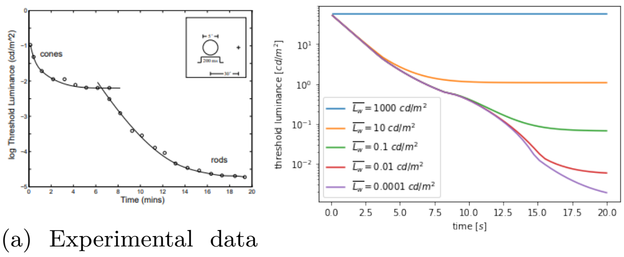

(a) Experimental data from Ferwerda et al.

(b) My model of dark adaptation

Figure 3: Threshold visibility of dark adaptation.

it does in fact show a similar shape when combined with my threshold function. Similarly, this function can simulate the photochemical model of light adaptation (Baker [1955] figs. 6 and 7) used for the initial state of dark adaptation if the user only briefly looks at a bright light. The new world adaptation luminance L_{wa} for the current frame is computed from the current optimal adaptation luminance $\overline{L_w}$ and the adaptation luminance L'_{wa} of the previous frame using a weight based on a constant parameter $0 < p < 1$ and the frame time delta dt seconds. The weight of the previous adaptation luminance decreases exponentially with increasing frame time.

$$L_{wa} = (1 - p^{dt}) \cdot \overline{L_w} + p^{dt} \cdot L'_{wa} \tag{7}$$

The parameter p can be tweaked for a faster or slower adaptation, and is chosen separately for light and dark adaptation. This function is evaluated for every frame to continually update the current simulated world observer adaptation as the user looks around. A plot of the threshold visibility function over the course of dark adaptation with $p = 0.5$ is given in fig. 3 along with Ferwerda et al.'s experimental data. The function is clearly non-linear in the mesopic region due to my choice of the k parameter in t_m. The different absolute threshold values during dark adaptation are likely caused by different experimental conditions between dark adaptation and fully adapted state. Because it is impractical to wait several minutes in VR for dark adaptation, especially if the user is standing, I decided to shorten the time course of dark adaptation from minutes to seconds similarly to Wernikowski et al..

However, as shown by Baker [1955, 1963], the photochemical model is insufficient to simulate light adaptation. Whereas in the photochemical model the threshold sensitivity slowly rises during light adaptation, physophysical data instead shows a sharp overcorrection in the threshold sensitivity followed by a slow decrease to the final threshold. Lacking access to the raw data, I decided to model this threshold function as

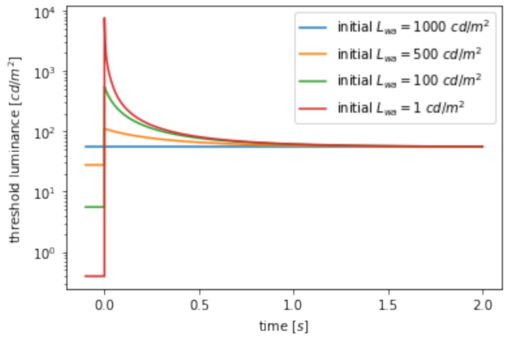

Figure 4: My model of threshold visibility of light adaptation.

Figure 5: Left: initial stages of light adaptation. Right: fully adapted to light.

follows:

$$t^*(L_{wa}, \overline{L_w}) = t_m(\overline{L_w}) \cdot \frac{t_m(\overline{L_w})}{t_m(L_{wa})} \tag{8}$$

After tweaking the parameter p used for the photochemical model, this function does in fact give the required shape for simulating light adaptation (fig. 4).

3.3 Tone reproduction

Ferwerda et al.'s tone reproduction operator is based on the approach developed by Ward [1994] using a threshold sensitivity function $t(L)$ to map world luminance L_w to display luminance L_d assuming a given light adaptation state L_{wa} for the simulated world observer and L_{da} for the display observer (i.e. user):

$$L_d(L_w) = \frac{t(L_{da})}{t(L_{wa})} \cdot L_w \tag{9}$$

For the display observer adaptation level I just use half of the approximate display luminance of the VR headset, which I set at 100 cd/m^2. This is in the photopic range of vision so I can simply use the threshold function for cones for the display observer: $t(L_{da}) = t_p(L_{da})$. Similarly, during dark adaptation I can just use the combined threshold function for the simulated world observer: $t(L_{wa}) = t_m(L_{wa})$.

As observed by Ferwerda et al., the increased threshold sensitivity during the initial stages of light adaptation when used with Ward's operator effectively dims the image when simulating the user being blinded by light. Because this does not match our daily experience, Ferwerda et al. proposed adding a constant offset to the displayed image. During light adaptation I therefore modify Ward's operator such that the display luminance mapped to $\overline{L_w}$ is kept constant:

Figure 6: Full (left) and simulated (right) colour discrimination for an images captured during a full moon.

$$L_d(L_w) = \frac{t_p(L_{da})}{t^*(L_{wa}, \overline{L_w})} \cdot (L_w - \overline{L_w}) + \frac{t_p(L_{da})}{t_m(\overline{L_w})} \cdot \overline{L_w} \quad (10)$$

An example of this effect can be seen in fig. 5.

3.4 Colour discrimination

Because our visual system is unable to discern colour at low ambient light levels, a faithful visual model must desaturate those parts of the image where colour discrimination is not possible. As I did not have access to the data used by Wernikowski et al. from Hunt [2004], I used their function with small changes:

$$\sigma = \begin{cases} 1 & \text{if } L_w > 3 \\ \frac{0.069}{0.069 + 1.409\, e^{-4.267\, L_w}} & \text{otherwise} \end{cases} \quad (11)$$

$$RGB_{display} = L_d(L_w) \cdot \left(\sigma \cdot \frac{RGB_{image}}{L_w} + (1 - \sigma) \right) \quad (12)$$

The function in Wernikowski et al. [2019] did not match their figure, likely due to typographical errors caused by missing minus symbols. Additionally, I inverted the meaning of the parameter such that $\sigma = 1$ produces a fully saturated colour and $\sigma = 0$ is fully desaturated for convenience. The effect of this operator can be seen in fig. 6.

4 Evaluation

The first method I tested, where the display parameters of the image are determined by the global luminance average of the image, is suitable for static images, but does not work well in VR. In many of the tested images it produced a “flat” look because it was adjusted to match neither the highlights nor the shadows. Some images (the night scenes in particular) looked fine, but switching between images felt a bit more jarring than with the other methods. This is a suitable method when the illumination of a scene can be constrained to a narrow range, but it is not suitable for viewing most HDR photographs.

The second tested method, where the display parameters are determined by the visible portion of the image, is better suited for displaying HDR photographs, but it feels very unnatural and quickly made me uncomfortable. This method feels similar to the “Auto Exposure” effect in Unity. The simulated adaptation time in Unity did prevent me from becoming uncomfortable, but neither effect felt natural when viewing HDR photographs in VR.

The final method simulating light and dark adaptation produced a comfortable viewing experience with the time parameters I chose for p in eq. (7). For light adaptation I chose $p = 0.1$, which matched my experience for small to medium changes in ambient luminance. The light adaptation was perhaps a little bit too fast for very large changes (night to clear day), but it did not seem overly unrealistic. For dark adaptation I chose the value $p = 0.5$ such that dark adaptation is mostly complete after 20 seconds. I found that this was a perhaps a little bit too fast for small changes, but not unreasonable. For large changes (daytime to nighttime) this was too fast to provide a realistic experience, but a slower change would be unbearable in VR.

I found that Wernikowski et al.’s colour discrimination operator worked well to desaturate images with luminance values in the mesopic range. For two of my images taken at night, I was not entirely happy with it, but I am unsure whether this is due to incorrect luminance calibration of the images or the colour discrimination operator. For a third night image (fig. 6), the colour discrimination operator provided a much more believable reproduction by desaturating the image compared to the raw values.

Despite being my preferred method, my simulation of light adaptation does have some limitations. I found that adding a constant offset to the whole image was not realistic in some situations, such as when looking directly at the sun. In this case, I would prefer to add a larger offset to an area surrounding the sun and a smaller offset to the rest of the image. Additionally, when looking at small bright objects, I found that the image produced a flickering appearance due to sampling variations when computing the average scene luminance. I was also unhappy with my weight function, in a few cases I found the final adaptation state too bright, others were too dim.

5 Conclusion

In this project I have proposed a simple model for simulating light and dark adaptation. This model provides a good viewing experience for HDR panoramic photographs in VR. It is possible to tweak the speed of light and dark adaptation based on the requirements of the application and preferences of the user. However, it does have some limitations.

Future work should consider a better method for determining the optimal final adaptation state to prevent flickering. A more complex method for simulating the “blinding” effect during early light adaptation may produce better results. More data is required to accurately model the time course of dark and especially light adaptation.

Additional resources

I have uploaded additional materials for this project to https://home.nulano.eu/nulano/AGIP/. This includes the source code of my model, a Windows executable of the project for the Oculus platform, the raw source images, and short video recordings of the light and dark adaptation simulation.

References

James A. Ferwerda, Sumanta N. Pattanaik, Peter Shirley, and Donald P. Greenberg. A model of visual adaptation for realistic image synthesis. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '96, page 249–258, New York, NY, USA, 1996. Association for Computing Machinery. ISBN 0897917464. doi: 10.1145/237170.237262. URL https://doi.org/10.1145/237170.237262.

Marek Wernikowski, Radosław Mantiuk, and Rafał Piórkowski. Preferred model of adaptation to dark for virtual reality headsets. In Ioannis Kompatsiaris, Benoit Huet, Vasileios Mezaris, Cathal Gurrin, Wen-Huang Cheng, and Stefanos Vrochidis, editors, *MultiMedia Modeling*, pages 118–129, Cham, 2019. Springer International Publishing. ISBN 978-3-030-05710-7. URL https://doi.org/10.1007/978-3-030-05710-7_10.

Howard D. Baker. Some direct comparisons between light and dark adaptation∗. *J. Opt. Soc. Am.*, 45 (10):839–844, Oct 1955. doi: 10.1364/JOSA.45.000839. URL https://opg.optica.org/abstract.cfm?URI=josa-45-10-839.

Howard D. Baker. Initial stages of dark and light adaptation∗†. *J. Opt. Soc. Am.*, 53 (1):98–103, Jan 1963. doi: 10.1364/JOSA.53.000098. URL https://opg.optica.org/abstract.cfm?URI=josa-53-1-98.

Greg Ward. Vii.2. - a contrast-based scalefactor for luminance display. In Paul S. Heckbert, editor, *Graphics Gems*, pages 415–421. Academic Press, 1994. ISBN 978-0-12-336156-1. doi: 10.1016/B978-0-12-336156-1.50054-9. URL https://www.sciencedirect.com/science/article/pii/B9780123361561500549.

R.W.G. Hunt. *The Reproduction of Colour*. John Wiley & Sons, Ltd, 2004. ISBN 9780470024270. doi: 10.1002/0470024275. URL https://onlinelibrary.wiley.com/doi/book/10.1002/0470024275.