Accurate and Visually Plausible Colour in Photorealistic Rendering

Billy Biggs, 97049311
Final project for CS788
University of Waterloo

Abstract

Photorealism in digitally rendered images requires the modeling of light transport through a scene, and also a modeling of how the eye views the incoming irradiance from the scene.

For modelling accurate light transport, this paper discusses the method of Stratified Wavelength Clustering (SWC), problems with this method, and implementation issues in both a path tracer and a bidirectional path tracer.

This paper also discusses issues with modeling the eye, such as glare, handling of high dynamic range scenes, chromatic adaption, and night-vision effects, and how this relates to a spectral renderer.

1 Introduction

The goal of a photorealistic rendering system is to store and present to a digital display device an image which appears ”realistic” to a human audience. This goal can be partitioned into three distinct elements:

1. A realistic model of light transport through a modeled scene of participants.
2. A realistic model of light perception by the human eye.
3. An inverse mapping from our model of the eye to intensity values suitable for CRT or LCD displays which give the desired effect to the viewer.

This paper discusses how performing colour operations on spectral samples rather than trichromatic units can improve accuracy for light transport problems, and how the stratified wavelength clustering approach is suitable for this problem. This paper then discusses how the output of a spectral raytracer can be incorporated into a model of light perception by the eye, and what colour processes are required to simulate the perception of a scene by a human viewer. For mapping onto the display, we discuss an interesting gamut mapping function.

2 Background

The first goal is to look at what colour representations are suitable for light transport operations both from an idealistic viewpoint and then from the implementation side. We first discuss spectral rendering, why it is necessary for accurate colour representation, and how it can be efficiently implemented. We then look at other issues which relate to perceptually accurate images, and how this can affect the design of our renderer.

2.1 The case for spectral rendering

The traditional RGB or XYZ basis functions are insufficient for modeling light transport through a scene [Peercy 1993, Raso and Fournier 1991]. Many effects depend on wavelength, and while some of these can be simulated using highly attenuated filters, they are insufficient for capturing differences in path direction due to wavelength, and also for representing interactions which involve spectra of multiple spikes.

Spectrum-dependent light transport effects we wish to simulate can be be divided into the following categories:

1. Specific absorbtion and transmission spectra defined as a function of wavelength.
2. Materials which shift power along the spectrum (fluorescence).
3. Colour dependence on the colour of the illuminant.
4. Changes in light path direction based on wavelength such as refraction and dispersion.

The first effect, modelling absorbtion and transmission spectra, can be simulated in some cases by storing a filter which can attenuate frequencies in accordance with their colour. For example, a yellow colour filter can be simulated in linear RGB space by storing an attenuation factor or power function under which to apply the multiplication operation. This can significantly change how colours react with each other, however, it is merely a heuristic and not an accurate description of colour interaction.

The effects of fluorescence can be modelled at each stage in RGB by converting to a spectral representation, applying the matrix transform, and converting back. However the conversion from RGB to a spectrum is highly unstable: there may be aliasing.

The third effect, colour dependence on the illuminant, is one way of describing how our perceived colour is insufficient for representing the actual colour. One hack for non-global illumination situations is to preprocess colours depending on the static light sources in the scene.

The final effect, modelling changes in light direction, is the most awkward to solve as most RGB colours are active all across the spectrum.

2.2 Conversion between spectral representations and standard basis functions

Human colour perception is limited to three dimensions corresponding to the three types of cones in the eye, and an additional sensitivity factor corresponding to the rods. A colour can also be represented as a spectrum, that is, as a continuous power distribution over the set of visible and near-visible wavelengths of light. Conversion between these two representations is performed using standardized matching functions. For example, from a power distribution $\Phi(\lambda)$, we can convert to the CIE XYZ colour space using the basis functions $\bar{x}$, $\bar{y}$, and $\bar{z}$ as follows:

$$X = k \int \Phi(\lambda)\bar{x}(\lambda)d\lambda$$

$$Y = k \int \Phi(\lambda)\bar{y}(\lambda)d\lambda$$

$$Z = k \int \Phi(\lambda)\bar{z}(\lambda)d\lambda$$

Where k is 683 lumens per Watt. This projects the infinite dimension space represented by the spectrum down to three dimensional XYZ space. The ability to reconstruct the original signal has been completely lost: that is, there are an infinite number of unique spectra which map to the same XYZ coordinates. This is the cause of metamers, that is, unique spectra which appear equivalent to the human eye.

The problem of metamers complicates conversion from RGB values into appropriate spectra. Colours for objects in the scene, as well as any texture data, must be converted to a spectral representation. The most straightforward method of performing this conversion is simply to evaluate some reasonable basis functions for the colours. For example, using the RGB basis functions $\bar{r}$, $\bar{g}$, and $\bar{b}$ as outputs from the RGB representation of the colour will definitely give a valid spectrum. However, these functions are not common in nature, and will lead to colour interactions that produce unnatural colours. In order to find a more pleasing metamer, one can assume that colour spectra for common objects and textures are not spiky, and so estimates should attempt to find a ”smooth” metamer for the desired colour.

2.3 Implementation options for spectral rendering

One method of approaching the spectral rendering problem is to apply more basis functions to spectral representations. While this does not address problems which require wavelength dependent changes in path direction, it allows for more accurate colours in many cases, and the cost is relatively small.

However, unless the basis functions used are completely independent box functions, then we still see colour error during operations such as addition or multiplication. If the boxes are too large, then we lose the ability to represent ”spikes” in the colour spectrum, which leads to poor colour depth in the image in general.

In a monte-carlo rendering setting, it makes the most sense to treat the colour spectrum as yet another integral to be estimated. This idea was described in more detail in [1] on Stratified Wavelength Clusters (SWC).

The basic idea is best described in terms of adding a new parameter to our path formulation integral: the wavelength of light. The final integral becomes:

$$Q_i = \int\int\int W_{ij}^Q(\lambda,\bar{x},\hat{\omega})L(\lambda,\bar{x},\hat{\omega})d\bar{x}d\hat{\omega}d\lambda$$

Where the weighting function for the intensity along the given path must include the conversion from the wavelength representation to the quantity Q required. That is:

$$W_{ij}^Q(\lambda,\bar{x},om\hat{e}ga) = k_Q\bar{q}(\lambda)W_ij(\bar{x},\hat{\omega})$$

Where $W_ij(\bar{x},\hat{\omega})$ is the weighting assigned to the given path (for when there are multiple samples for paths of the same length, as in bidirectional path tracing).

In order to obtain an estimate of this integral, we apply monte-carlo sampling, both on the the path lengths and on the wavelength component. To do this, we need a distribution function for the wavelengths of light, and then some way of gathering enough samples to make our estimate worthwhile.

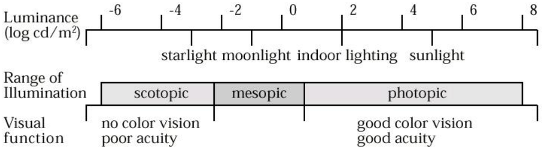

Figure 1: Visual modes and their associated luminances, after Hood 1986

2.4 Sampling the wavelength distribution from the light source

The SWC approach involves choosing a light source at random and then choosing a set of wavelengths sampled from its power distribution, stratified over the cummulative distribution. Given a light source, we choose K random values from K stratified intervals in $[0 - 1]$ and select frequencies of the light source according to its normalized power output.

The first issue with this formulation occurs when the there are many light sources in the scene with largely different power spectrums. A possible solution here is to weight the probability of each light source based on its total output power, and then include this as part of the probability for each wavelength.

Another issue with this formulation is handling of flourescence. In a bidirectional path tracing situation, this may result in different clusters of wavelengths on different path permulations. This problem will not occur in a traditional path tracing situation, since each sample is constructed from a single path.

One possible solution is to create the light path, and then choose K uniform random wavelength values from the resulting set of unique wavelengths, and strip the rest of the wavelength results from the cluster at each node. Then this new cluster can be used when evaluating the eye path, and it is guarenteed to be valid at all points along the light path. While the existance of fluorescent objects in the scene has already increased our error, by choosing K random wavelengths we avoid adding bias for frequencies modified by the fluorescent transform, while still allowing these new frequencies to occur.

The final problem with SWC in general is the degredation of the estimate when refraction is encountered. At this point, we lose the benefit of transporting multiple wavelengths per sample, and the problem degrades to a path per wavelength solution.

2.5 Output of a spectral raytracer for tone and hue mapping

A spectral raytracer is most useful for input into a tone mapping system. While the eye’s colour representation is eventually trichromatic, there are situations which require more knowledge from the original spectral representations of the image.

The first situation is gaining an accurate measure of the scotopic receptor, the rods in the eye. There are three main modes of activation of the light receptors in the eye: Scotopic, Mesopic, and Photopic. See **figure 1** for a graphic description of the luminance levels associated with different visual modes. Photopic vision corresponds to daylight levels, while scotopic vision occurs at low levels of light.

The scotopic luminous efficiency function is identical to the photopic luminous efficiency function, except that it is shifted down-

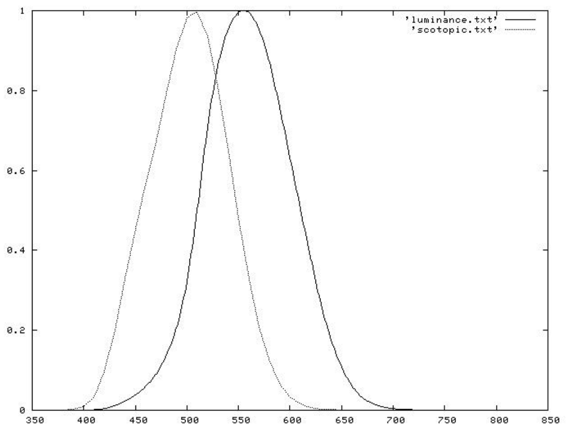

Figure 2: Plot of $V'(\lambda)$, the scotopic luminous efficiency function, and $V(\lambda)$, the photopic luminous efficiency function, from [6].

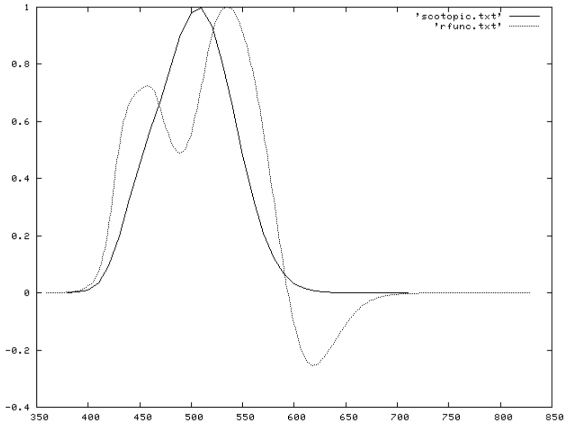

Figure 3: Comparison of the R function given by [3] versus the target CIE $V'(\lambda)$ function.

wards along the spectrum. See **figure 2** for a plot of the photopic and scotopic luminous efficiency functions.

From [3], the rod signal can be estimated from only the XYZ sources as follows:

$$R = -0.702X + 1.039Y + 0.433Z$$

However this has problems: it is a poor linear approximation and negative values must be dealt with by clamping. A plot of this function versus the CIE $V'(\lambda)$ function is given in figure 3. For more accuracy in both colour reproduction and night-vision effects, it's much easier to apply the scotopic matching function directly to the spectrum along with the matching functions for X, Y, and Z.

2.6 Modeling glare

Effects of glare, that is, objects with high brightness compared to their surroundings tend to flare around them.

The effects of glare occur due to scattering as light enters the eye, and are not modeled in terms of rod and cone interactions. However, the visual effects of glare are dependent on the visual function of the observer, that is, the lenticular halo is much more pronounced when the eye is in the scotopic state [4].

Unfortunately, the glare mapping not only requires knowledge of the visual mode, it also must estimate the actual intensity of the display device in order to estimate what intensities will appear as real glare and not need additional processing. Because of this, the above transformations for visual mode and the difference between the maximum intensity of the image and that of the display device must be calculated.

2.7 Tone Mapping: Chromatic Adaption and Dynamic Range Compression

Two visual effects which improve quality of generated images are chromatic adaption and dynamic range compression. Chromatic adaptation occurs when a coloured illuminant appears white after the eye has adjusted to it. That is, eventually a room lit by a red-tinted light will appear fully coloured. For handling of high dynamic range scenes, the eye adjusts to the change in a non-linear fashion,

There have been many attempts at determining appropriate operators to conpensate for high dynamic range in rendered scenes. One modern approach that is representative of modern tone mapping ideas is described in [3]. The process involves taking an image in terms of the rod component (V'), L, M and S components (for long, medium and short sensitivity cones), and performing a transform on the image in 6 bandpass filtered representations, corresponding to decreasing levels of detail. At each contrast level, a logarithmic transform is applied to effectively perform a compression on the dynamic range separately on each level of the detail pyramid. Chromatic adaptation is performed in a similar fashion, by looking at the achromatic signal separately from the chromatic signals and again performing a non-linear transform. The results are then combined to a final 3-component image.

2.8 Modeling 'night-vision'

It was noted in [2] that current models for tone mapping fail to demonstrate the 'blue' appearence of scotopic scenes. To compensate for this, they suggested a simple method of hue mapping to increase the blue in an image, based on the original samples in XYZ, scotopic sensitivity V', plus an additional parameter s which measures how scotopically sensitive the viewer is.

There is no quantitative data to properly measure this effect yet available, and so therefore the transformation of making the image more 'blue', while physically based, is not known to be accurate. Therefore, its use is a tradeoff between ultra-accuracy and plausible images.

The computation is defined on the XYZ representation, a chosen chromaticity for the blue we wish to shift to x_b and y_b (they advocate a blue-point of $(0.25, 0.25)$), as well as the calculated chromaticities x and y as follows:

$$x = (1-s)x_b + sx$$
$$y = (1-s)y_b + sy$$
$$Y = 0.4468(1-s)V' + sY$$
$$X = xY/y$$
$$Z = X/x - X - Y$$

The s value used should be the same as is used to calculate the glare option mentioned above.

The process of hue mapping can be performed after the image has been converted from tone and dynamic range mapping. The result will be a blurred image due to the decreased sensitivity of the eye in low-light situations, and will then be blue-shifted.

3 The processing pipeline

Given all of the above goals and background, we can now describe how an implementation can achieve the goals of visual accuracy with a tractable amount of complexity.

The proposed chain of events is:

1. Convert input colours to spectral representations.

   Since we have chosen to model light transport in the spectral domain, any input colours or filters defined in a tri-chromatic system must be converted to spectral representations. This usually involves determining a mapping based on the colour and a smoothness value.

2. Render the scene using SWC to four components: X,Y,Z,V'.

   As described above, the SWC method is applied to the scene to output four channels, the three usual colour component representations, plus an additional representation of the scotopic response function V', which we will use to calculate the night-vision hue shift, as well as take part in our model of chromatic adaption.

3. Compute visual function based on output intensity.

   In order to perform visual processing, we must determine under which visual mode the scene will be viewed, scotopic, photopic, or "in-between" (mesopic). This involves measuring the luminous intensity (or radiant intensity) of the image and choosing a mode based on the maximum value for the log of the scene luminance (calculation of our s value).

4. Apply a digital glare filter on the final XYZV' image.

   The glare filter described in [4] depends on both the visual mode (scotopic, photopic, etc) and also the maximum representable intensity. However, [3] claims that this conversion must be performed before tone mapping.

5. Filter and convolve the rod, S, M, L representations of the image for chromatic adaption and dynamic range compression.

   This involves filtering into multiple bandpassed segments at different spacial frequencies, applying a non-linear transform, and converting back.

6. Convert from our four-channel representation back to three-channel colours, applying a hue shift for the rods depending on scotopic sensitivity.

   For modeling the night-vision effect, we apply the mapping as described in [2], by measuring the scotopic adaption and potentially applying a heavy shift towards the blue area of the scene.

7. Convert from our final XYZ image to in-gamut and gamma corrected R'G'B' for output to a digital display device.

4 Implementation results

The first process that was completed was the addition of a spectral representation for colour and functions to map between this representation and CIE XYZ. As an approximation, input RGB values were converted to a spectral representation simply by using the curve defined by the RGB primaries of the input colours. This favours representations of pure red, green, and blue, but outputs unlikely spectra for in-between colours, and is an unreasonable approach in general.

The SWC method of sampling paths using their spectra was then applied to the path tracer. As scene in the provided images, the resulting effects showed an increase in the overall colour accuracy. Bidirectional methods were applied (as per [5]) to the path tracing algorithm to increase the efficiency for generation of less noisy scenes.

The output was evaluated in terms of X,Y,Z,V' representations, and the calculations performed for estimating the scotopic sensitivity of the eye.

5 Conclusions and future work

This paper describes in detail the process of generating a reasonable and photorealistic image in a tractable amount of computation time. We describe how stratified wavelength clustering is a reasonable approach for modeling spectral-based colour interactions, and how to convert the output to formats which are reasonable for input to a tone mapper. We describe the full tone mapping process and what options are available for creating a more realistic image.

References

[1] Glenn F. Evans and Michael D. McCool. Stratified wavelength clusters for efficient spectral monte carlo rendering. *Graphics Interface*, pages 42–49, June 1999.

[2] Henrik Wann Jensen, Fredo Durand, Michael M. Stark, Simon Premoze, Julie Dorsey, and Peter Shirley. A physically-based night sky model. *SIGGRAPH Proceedings*, 2001.

[3] Sumanta N. Pattanaik, James A. Ferwerda, Mark D. Fairchild, and Donald P. Greenberg. A multiscale model of adaptation and spatial vision for realistic image display. *SIGGRAPH Proceedings*, pages 287–298, August 1998.

[4] Greg Spencer, Peter Shirley, Kurt Zimmerman, and Donald P. Greenberg. Physically-based glare effects for digital images. *SIGGRAPH Proceedings*, pages 325–334, November 1995.

[5] Eric Veach. Robust monte carlo methods for light transport simulation. 1997.

[6] G. Wyszecki and W.S. Stiles. Color science: Concepts and methods, quantitative data and formulae (2nd edition). 1982.