

Tone Reproduction and Physically Based Spectral Rendering

Kate Devlin¹ Alan Chalmers¹ Alexander Wilkie² Werner Purgathofer²

1 – Department of Computer Science
University of Bristol

2 – Institute of Computer Graphics and Algorithms
Vienna University of Technology

Abstract

The ultimate aim of realistic graphics is the creation of images that provoke the same responses that a viewer would have to a real scene. This STAR addresses two related key problem areas in this effort which are located at opposite ends of the rendering pipeline, namely the data structures used to describe light during the actual rendering process, and the issue of displaying such radiant intensities in a meaningful way.

The interest in the first of these subproblems stems from the fact that it is common industry practice to use RGB colour values to describe light intensity and surface reflectancy. While viable in the context of methods that do not strive to achieve true realism, this approach has to be replaced by more physically accurate techniques if a prediction of nature is intended.

The second subproblem is that while research into ways of rendering images provides us with better and faster methods, we do not necessarily see their full effect due to limitations of the display hardware. The low dynamic range of a standard computer monitor requires some form of mapping to produce images that are perceptually accurate. Tone reproduction operators attempt to replicate the effect of real-world luminance intensities.

This STAR report will review the work to date on spectral rendering and tone reproduction techniques. It will include an investigation into the need for spectral imagery synthesis methods and accurate tone reproduction, and a discussion of major approaches to physically correct rendering and key tone mapping algorithms. The future of both spectral rendering and tone reproduction techniques will be considered, together with the implications of advances in display hardware.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Viewing Algorithms I.3.7 [Computer Graphics]: Color, shading, shadowing, and texture

1. Introduction

The ultimate aim of realistic graphics is the creation of images that provoke the same response and sensation as a viewer would have to a real scene, i.e. the images are physically or perceptually accurate when compared to reality. This requires significant effort to achieve, and one of the key properties of this problem is that the overall performance of a photorealistic rendering system is only as good as its worst component.

In the field of computer graphics, the actual image synthesis algorithms – from scanline techniques to global illumination methods – are constantly being reviewed and improved, but two equally important research areas at opposite

ends of the rendering pipeline have been neglected by comparison: the question which entities are used in rendering programs to describe light intensity during the calculations performed by the rendering algorithms, and the mapping of the luminances computed by these algorithms to the display device of choice. Weaknesses of a system in both areas can make any improvements in the underlying rendering algorithm totally pointless. Consequently, good care has to be taken when designing a system for image synthesis to strike a good balance between the capabilities of the various stages in the rendering pipeline.

In this paper we review the state of the art on these two topics in an interleaving manner. We first present the basic

problems in both areas in sections 2 and 3, and then discuss previous work on tone mapping in section 4 and for spectral rendering in section 5.

2. Tone mapping

While research into ways of creating images provides us with better and faster methods, we usually do not see the full effect of these techniques due to display limitations. For accurate image analysis and comparison with reality, the display image must bear as close a resemblance to the original image as possible. In situations where predictive imaging is required, tone reproduction is of great importance to ensure that the conclusions drawn from a simulation are correct (Figure 1).

2.1. The need for accurate tone reproduction

Tone reproduction is necessary for two main reasons: the first is to ensure that the wide range of light in a real world scene is conveyed on a display with limited capabilities; and the second is to produce an image which provokes the same responses as someone would have when viewing the scene in the real world. Physical accuracy alone of a rendered image does not ensure that the scene in question will have a realistic visual appearance when it is displayed. This is due to the shortcomings of standard display devices, which can only reproduce a range of luminance of about 100:1 candelas per square metre (cd/m^2), as opposed to human vision which ranges from 100 000 000:1, from bright sunlight down to starlight, and an observer's adaptation to their surroundings also needs to be taken into account. It is this high dynamic range (HDR) of human vision that needs to be scaled in some way to fit a low dynamic range display device.

In dark scenes our visual acuity — the ability to resolve spatial detail — is low and colours cannot be distinguished. This is due to the two different types of photoreceptor in the eye: rods and cones. It is the rods that provide us with achromatic vision at these *scotopic* levels, functioning within a range of 10^{-6} to 10 cd/m^2 . Visual adaptation from light to dark is known as *dark adaptation*, and can last for tens of minutes; for example, the length of time it takes the eye to adapt at night when the light is switched off. Conversely, *light adaptation*, from dark to light, can take only seconds, such as leaving a dimly lit room and stepping into bright sunlight. The cones are active at these *photopic* levels of illumination, covering a range of 0.01 to 10^8 cd/m^2 . The overlap (the *mesopic* levels), when both rods and cones are functioning, lies between 0.01 to 10 cd/m^2 . The range normally used by the majority of electronic display devices (cathode ray tubes, or CRTs) spans from 1 to 100 cd/m^2 . More detailed information on visual responses with regard to tone reproduction can be found in the papers by Ferwerda et al., Pattanaik et al. and Tumblin ^{13, 48, 50, 81}.

Despite a wealth of psychophysical research, our knowl-

edge of the Human Visual System (HVS) is still limited, but its ability to perceive such a wide dynamic range in the real-world requires some form of reproduction to produce perceptually accurate images on display devices. Changes in the perception of colour and of apparent contrast also come into play when mapping values to a display device. The development of new psychophysically-based visual models seeks to address these factors. To date, methods of tone mapping tend to concentrate on singular aspects for singular purposes. This approach is understandable given the deficit in HVS knowledge, but is inefficient as the HVS responds as a whole, rather than as isolated functions. New psychophysical research is needed to address the workings of the HVS in their totality.

2.2. Tone mapping: art, television and photography

Tone mapping was developed for use in television and photography, but its origins can be seen in the field of art where artists make use of a limited palette to depict high contrast scenes. It takes advantage of the fact that the HVS has a greater sensitivity to relative rather than absolute luminance levels ²⁶. Initial work on accurate tone reproduction was insufficient for high dynamic range scenes. Either the average real-world luminance was mapped to the display average, or the maximum non-light source luminance was mapped to the maximum displayable value. However, the process failed to preserve visibility in high dynamic range scenes as the very bright and very dimmed values were clamped to fall within the display range. Also, all images were mapped irrespective of absolute value, resulting in the loss of an overall impression of brightness ³⁸.

The extensive use of tone reproduction in photography and television today is explained in Hunt's "The Reproduction of Colour in Photography, Print and Film" ²⁶ and Poynton's "A Technical Introduction to Digital Video" ⁵⁵, which give comprehensive explanations on the subject. This research area is outside the scope of this paper, and it is suggested that readers with an interest in this area refer initially to these works.

2.3. Gamma correction

Gamma is a mathematical curve that represents the brightness and contrast of an image. Brightness is a subjective measurement, formally defined as "the attribute of a visual sensation according to which an area appears to emit more or less light" ⁵⁴. It describes the non-linear tonal response of the display device and compensates for the non-linearities. For CRTs, the use of RGB values to express colour is actually specifying the voltage that will be applied to each electron gun. The luminance generated is not linearly related to this voltage. In actuality, luminance produced on the display device is approximately proportional to the applied voltage raised to a power of 2.5, although the actual value of the




Figure 1: Ideal tone reproduction process

exponent varies⁵⁴. Gamma correction seeks to rectify these anomalies, but a gamma of 1.0 (where viewing conditions between the original scene luminance and reproduced luminance are identical) is not always desirable. Psychophysical studies have shown that a higher gamma value is preferable for all but the brightest conditions⁸¹. For CRTs, an inverse power function of 1/2.2 is applied to the RGB data before display. Although this goes some way towards correcting the data, there is still scope for variation. Most monitors provide brightness and contrast controls. Correction may also have been applied to the image data or in the user software. These potential areas for correction can lead to inconsistencies and it cannot be assumed that an approximation of an ideal display has been achieved.

When displaying images that have been modified by a tone reproduction operator, a gamma corrected monitor is desired.

3. Spectral Rendering

For the purposes of truly predictive photorealistic rendering it is essential that no effect which contributes to the interaction of light with a scene is neglected. Most aspects of object appearance can be accounted for by using just the laws of geometric optics, comparatively simple colour space descriptions of surface reflectivity, tristimulus representations of colour and light, and can nowadays be computed very efficiently through a variety of common rendering algorithms.

However, due to the approximation inherent in performing rendering in colour space, several important physical effects, namely fluorescence, diffraction, dispersion and polarization, are still rarely – if at all – supported by contemporary image synthesis software. Partly as a consequence of this, most available rendering software cannot be used to reliably predict illumination intensities in scenes with nontrivial geometries and surface reflectancies, which in turn precludes the ability to truly match the visual impression of a human observer with a virtual scene.

Although for instance the pioneering graphics group at Cornell University already noted early in the 1980ies that colour computations in a renderer have to be performed in spectral space if the output is to be used for predictive

purposes²⁴, the computer graphics mainstream has up to now avoided spectral rendering techniques in favour of the seemingly more robust and less complicated colour space approach; occasional publications and surveys²³ on the topic have not led to a breakthrough in acceptance yet.

In this paper, we use the term “spectral rendering” to mean image synthesis methods which use some kind of representation of the associated light spectrum for colour values (i.e. light intensities), as opposed to conventional systems that perform these calculations with tristimulus colour values.

Most standard computer graphics textbooks do not go into detail concerning spectral rendering (if they mention the problem at all); notable exceptions in this respect are Hall²², Glassner¹⁹ and Shirley⁶⁷.

In section 3.2 we aim to give an overview over the effects which are only tractable using such systems, and discuss implementation issues along with an overview of spectral rendering systems in section 5.

3.1. Colour and Light

Since the ultimate goal of the image synthesis step in a realistic rendering pipeline is the computation of the colour associated with a given pixel, we have to briefly discuss the connection between light intensities and colour values at this point. Readers interested in a more thorough treatise are for instance referred to the excellent anthology by Nassau⁴³, which goes into great detail with respect to many areas which are usually omitted from computer graphics literature. For the definitive reference work on this topic, the reader is referred to Wyszecki and Stiles⁹².

The human eye is sensitive to electromagnetic radiation in just a tiny segment of the spectrum, namely from about 380nm to 780nm; individual sensitivities can vary considerably between subjects, and are also dependent on the physical state of the observer.

However, the human eye is not a full spectral sensor, but rather has four types of receptor cells, which are responsible for luminosity and overlapping regions in the short, medium and long wavelengths, respectively. Response curves of the

human eye for three selected wavelengths (red, green and blue) were determined through colour matching experiments, and standardized by the CIE in 1932⁹² for a particular set of viewing conditions. These RGB colour matching curves were normalized to yield a device-independent colour space, CIE XYZ. Since then numerous detail improvements to the measurement process and additions for different viewing conditions have been made, but the original CIE XYZ colour space is still the standard for device-independent colour description.

3.1.1. Conversions

The process of converting a given spectral power distribution to its corresponding CIE XYZ colour value is straightforward and described in detail in many computer graphics textbooks^{22,43,19,67}. It basically just requires convolution of the given spectral power distribution S by the appropriate matching function x, y, z for the channel in question:

$$X = \int_{380}^{780} x(\lambda)S(\lambda)d\lambda \quad (1)$$

$$Y = \int_{380}^{780} y(\lambda)S(\lambda)d\lambda \quad (2)$$

$$Z = \int_{380}^{780} z(\lambda)S(\lambda)d\lambda \quad (3)$$

The resulting XYZ value can be converted to a RGB value through multiplication by a transformation matrix which is unique for the colour space of each RGB output device; if the resulting colour has negative components and hence is outside the display gamut, a gamut reduction technique of some kind has to be applied.

The reverse transformation is not an unique operation – infinitely many spectra correspond to a given RGB triplet – and is therefore fraught with difficulties, since not any of these spectra, but rather the metamer best suited for the task at hand has to be found.

Methods for deriving such a spectral distribution for a given RGB value have been proposed by Glassner¹⁸ and more recently by Smits⁷¹; the latter is an improvement insofar as it actively searches metamer space for a spectrum which is physically plausible.

3.1.2. Representations

If a rendering system is to use spectral distributions for its light intensity calculations, there are several options for the storage of these functions, which are usually quite smooth, but can – for example in the case of fluorescent illuminants – also feature sharp, high-frequency spikes. The obvious trade-off here is between accuracy and computation speed; while no-one disputes that sampling a given power spectrum at 5nm or even 1nm intervals will yield satisfactory results, the large memory requirements and convolution times

of such approaches make them impractical even on modern systems.

The important techniques in practical use are sparse direct sampling – usually at intervals above 10nm – and basis function approaches⁵¹, which are sometimes also referred to as *linear methods*, since they linearly combine a set of given basis functions.

The former suffer from lack of accuracy if fluorescent light sources with pronounced spikes are to be represented with just a few samples, and the latter have the problem that a given small set of basis functions is usually just suitable for a given set of input spectra, but not for arbitrary power distributions.

On the other hand, if large numbers of basis functions are used the advantages of this approach – a lower number of coefficients is needed – is eroded to the point where the simpler direct sampling method is more efficient at lower computational cost.

Raso et al.⁵⁹ and later Geist et al.¹⁵ proposed to use polynomials to represent spectra. Both this and the basis function approach are quite compact, but suffer from the fact that spectral multiplications are of the order $O(n^2)$; directly sampled spectra just require $O(n)$ for this operation.

Adaptive techniques that aim to represent just the pertinent parts of the visible spectrum have been proposed by Deville et al.⁹ and Rougeron et al.⁶²; however, neither of these has been tested in a production renderer so far, and it therefore remains to be seen whether the gains in efficiency are worth the additional complexity of such techniques.

Recently, an efficient hybrid or composite approach to spectral sampling has been proposed by Sun et al.⁷⁸. Low-order basis functions are used for representation of the overall spectral shape, and spectral spikes are maintained separately. In this way a very high accuracy is maintained even though just a few coefficients have to be stored.

A possible drawback of this approach are the comparatively complex spectral multiplications. This problem is common to all more sophisticated spectral representation techniques: on modern processors a simple multiplication loop – as required by the approach of spectra directly sampled at similar intervals – can be vectorized, while complex folding operations take much longer to compute.

3.2. Effects for which Spectral Rendering is a Necessity

In realistic graphics, there are two categories of effects which require a spectral renderer: physical phenomena which cannot be computed accurately unless more than tristimulus values are used (although approximations are possible in some cases), and what ultimately are perception issues which cannot be resolved by computations in colour space.

The latter group would include the problem of

metamerism, and issues related to perceptually accurate tone reproduction operators which might require spectral input data in order to accurately mimic the behaviour of the human eye ⁴⁹.

In this section, we will discuss the former group – the physical effects – in more detail.

3.3. Dispersion in Dielectric Materials

Dispersion occurs where polychromatic light is split into its spectral components on a refractive material boundary due to the fact that the index of refraction in transparent materials is dependent on the wavelength of the incident light. Usually this dependency on wavelength is non-linear and related to material constants that have to be measured in experiments.

The perceived result of this effect are usually coloured fringes in glass objects, and rainbow caustics cast by prisms or crystals; several researchers have investigated these phenomena in the past ^{80, 94, 5, 89}.

While wavelength dependent refraction is theoretically also possible in a tristimulus based renderer, the low number of independent colour bands prohibits a faithful representation of the resulting rainbow effects in such a system.




Figure 2: Split of an incident white light beam into its spectral components in a prism.

A topic which is generally not covered in computer graphics textbooks (with the notable exception of Glassner ¹⁹) is how the wavelength-dependency of the IOR can be described analytically.

The most widely used method of specifying the dispersion curve for materials in the visual range is to use the so-called *Sellmeier approximation* ^{3, 19}. Several basically similar forms exist that differ only in the number of empirical constants in structurally similar equations. The number of these constants usually depends on the measurement process by which the data for the approximation is obtained and the




Figure 3: Refractive indices for some dielectric materials. From top: diamond, lead crystal and several normal glass types. Notice the varying amount of dispersion and non-linearity for different materials.

associated desired accuracy, and is specific to the source of the data.

A typical example is the glass catalog of the company Schott Glaswerke ⁶⁶, which is one of the worldwide leading suppliers of technical glass. In the catalog the technical data of the several hundred types of glass that the company sells is listed, and for specifying dispersion the form

$$n^2(\lambda) - 1 = \frac{B_1\lambda^2}{\lambda^2 - C_1} + \frac{B_2\lambda^2}{\lambda^2 - C_2} + \frac{B_3\lambda^2}{\lambda^2 - C_3} \quad (4)$$

based on three resonance frequencies is used, where n is the index of refraction at wavelength λ .

The catalog lists coefficient values of B_n and C_n for the different glass types (ranging from normal window glass to highly dispersive lead crystal). In this particular case one can compute the index of refraction for wavelengths from ultraviolet to far infrared with a relative error of less than $1.0E-5$ from just six coefficients per glass type. This makes the catalog a valuable source for accurate dispersion data, especially since it can be downloaded from the company website free of charge and contains specimens of all the main basic glass types (i.e. flints, crowns, lead crystal aso.).

There are also other sources of similar freely available material measurements where one can obtain measurements of dielectric materials other than glass (e.g. diamond), both on the web and more importantly in book form ⁴⁶.

3.4. Polarization

Polarization of light has received particularly little attention in the rendering community because – while of course being essential for specially contrived setups that for instance contain polarizing filters – it seemingly does not contribute very prominent effects to the appearance of an average scene.

This misconception is in part fostered by the fact that the human eye is normally not thought of as being capable of distinguishing polarized from unpolarized light. However, contrary to common belief trained observers can identify strongly polarized light with the naked eye.

Named after its discoverer, the effect is known as *Haidinger's brush* and is described by Minnaert in his book about light in outdoor surroundings⁴⁰. It is for instance readily manifest to any user of an LCD monitor who knows what to look for⁵³; once one is aware of the typical two-coloured pattern induced by polarized light, one frequently spots occurrences in everyday life.

One of the main areas where polarization in fact does make a substantial difference to the overall radiance distribution are outdoor scenes; this is due to the – under certain circumstances quite strong – polarization of skylight on clear days, as one can find documented in G. P. Können's book³³ about polarized light in nature. But since outdoor scenes are currently still problematical for photorealistic renderers for a number of other, more obvious reasons (e.g. scene complexity and related global illumination issues), this has not been given a lot of attention yet.

Also, although comparatively sophisticated analytical skylight models which are even partially based on spectral radiance measurements have been presented recently⁵⁶, no mathematical description of the polarization patterns found in a clear sky – as shown e.g. in figure 4 – has been presented so far. Other known effects which depend on polarization




Figure 4: Skylight polarization levels in percent for two different solar elevations. *S* solar position, *A* Arago point, *Br* Brewster point, *Ba* Babinet point. Redrawn from Mütze et al.⁴¹

support are darkening or discolourization patterns in metal objects and their specular reflections in dielectric materials, the darkening of certain facets in transparent objects such as crystals, and a large number of scattering phenomena in air and liquids³³.

3.4.1. Causes of Light Polarization

Apart from skylight, it is comparatively rare for light to be emitted in polarized form. In most cases, polarized light is the result of interaction with transmitting media or surfaces.

The correct simulation of such processes is at the core of predictive rendering, so a short overview of this topic recommends itself.

The simplest case is that of light interacting with an optically smooth surface. This scenario can be adequately described by the *Fresnel equations*, which are solutions to Maxwell's wave equations for light wavefronts. They have been used in computer graphics at least since Cook and Torrance proposed their reflectance model⁸, and most applications use them in a form which is simplified in one way or another.




Figure 5: Four examples of the patterns traced out by the tip of the electric field vector in the X-Y plane: a) shows light which is linearly polarized in the vertical direction; the horizontal component E_x is always zero. b) is a more general version of linear polarization where the axis of polarization is tilted by an angle of α from horizontal, and c) shows right circular polarized light. The fourth example d) shows elliptically polarized light, which is the general case of equation (5). (Image redrawn from Shumaker⁶⁸)

3.5. Polarized Light

While for a large number of purposes it is sufficient to describe light as an electromagnetic wave of a certain frequency that travels linearly through space as a discrete ray (or a set of such rays), closer experimental examination reveals that such a wavetrain also oscillates in a plane perpendicular to its propagation. The exact description of this phenomenon requires more than just the notion of radiant intensity, which the conventional representation of light provides.

The nature of this oscillation can be seen from the microscopic description of polarization, for which we closely follow that given by Shumaker⁶⁸.

We consider a single steadily radiating oscillator (the light source) at a distant point of the negative Z-axis, and imagine that we can record the electric field present at the origin due to this oscillator. Note that the electric and magnetic field vectors are perpendicular to each other and to the propagation of the radiation, and that the discussion could equally

well be based on the magnetic field; which of the two is used is not important.

Except at distances from the light source of a few wavelengths or less, the Z component of the electric field will be negligible and the field will lie in the X–Y plane. The X and Y field components will be of the form

$$\begin{aligned} E_x &= V_x \cdot \cos(2\pi \cdot \mathbf{v} \cdot \mathbf{t} + \delta_x) \quad [\text{V} \cdot \text{m}^{-1}] \\ E_y &= V_y \cdot \cos(2\pi \cdot \mathbf{v} \cdot \mathbf{t} + \delta_y) \end{aligned} \quad (5)$$

where V_x and V_y are the amplitudes [$\text{V} \cdot \text{m}^{-1}$], \mathbf{v} is the frequency [Hz], δ_x and δ_y are the phases [rad] of the electromagnetic wavetrain, and t is the time [s]. Figure 5 illustrates how this electric field vector E changes over time for four typical configurations.




Figure 6: Geometry of a ray–surface intersection with an optically smooth phase boundary between two substances, as described by the equation set (6). A transmitted ray T only occurs in when two dielectric media interface; in this case, all energy that is not reflected is refracted, i.e. $T = I - R$. The E -vectors for the transmitted ray $E_{t\parallel}$ and $E_{t\perp}$ have been omitted for better picture clarity. The $(E_{\parallel}, E_{\perp})$ components here correspond to the (x, y) components in the drawing on the left.

3.5.1. Fresnel Terms

In their full form (the derivation of which can e.g. be found in ⁶⁹), they consist of *two* pairs of equations, of which only the first is usually quoted in computer graphics literature. According to the reflection geometry in figure 5, the first pair determines the proportion of incident light which is reflected separately for the x and y components of the incident wavetrain. This relationship is commonly known, and can be found in numerous computer graphics textbooks.

The second pair, which is much harder to find ⁹¹, describes

the *retardance* that the incident light is subjected to, which is the relative phase shift that the vertical and horizontal components of the wavetrain undergo during reflection. In figure 7 we show the results for two typical materials: one conductor, a class of materials which has a complex index of refraction and is always opaque, and one dielectric, which in pure form is usually transparent, and has a real–valued index of refraction.

We quote the Fresnel equations for a dielectric–complex interface. This is the general case, since only one of two media at an interface can be conductive (and hence opaque), and a dielectric–dielectric interface with two real–valued indices of refraction can also be described by this formalism.

$$\begin{aligned} F_{\perp}(\theta, \eta) &= \frac{a^2 + b^2 - 2a \cos \theta + \cos^2 \theta}{a^2 + b^2 + 2a \cos \theta + \cos^2 \theta} \\ F_{\parallel}(\theta, \eta) &= \frac{a^2 + b^2 - 2a \sin \theta \tan \theta + \sin^2 \theta \tan^2 \theta}{a^2 + b^2 + 2a \sin \theta \tan \theta + \sin^2 \theta \tan^2 \theta} F_{\perp}(\theta, \eta) \\ \tan \delta_{\perp} &= \frac{2 \cos \theta}{\cos^2 \theta - a^2 - b^2} \\ \tan \delta_{\parallel} &= \frac{2b \cos \theta [(n^2 - k^2)b - 2nka]}{(n^2 + k^2)^2 \cos^2 \theta - a^2 - b^2} \\ &\text{with} \\ &\eta = n + ik \quad (\text{the complex IOR}) \\ 2a^2 &= \sqrt{(n^2 - k^2 - \sin^2 \theta)^2 + 4n^2 k^2} + n^2 - k^2 - \sin^2 \theta \\ 2b^2 &= \sqrt{(n^2 - k^2 - \sin^2 \theta)^2 + 4n^2 k^2} - n^2 + k^2 + \sin^2 \theta \end{aligned} \quad (6)$$

F_{\parallel} is the reflectance component parallel to the plane of incidence, and F_{\perp} that normal to it. Under the assumption that one is only interested in the radiant intensity of the reflected light, this can be simplified to the commonly used average reflectance $F_{\text{average}} = (F_{\perp} + F_{\parallel})/2$. δ_{\perp} and δ_{\parallel} are the retardance factors of the two wavetrain components.

3.6. Fluorescence

While the polarization of light at a phase boundary is a comparatively macroscopic phenomenon, fluorescence is caused by processes within the molecules that are responsible for the colour of an object. The key point is that re–emission of photons that interact with matter does not necessarily occur at the same energy level – which corresponds to a certain frequency and ultimately colour – at which they entered ^{19, 58}.

Both the case of re–emission at lower energy levels and the case of two lower energy photons being “combined” into a single higher–energy photon are common. However, for the purposes of computer graphics only the first case is of major interest, since it is the governing phenomenon behind fluorescent pigments such as for instance Dayglo® paint.




Figure 7: Fresnel reflectivities F_{\parallel} , F_{\perp} and F_{average} (dashed lines), as well as parallel and perpendicular retardance values for copper (red) and lead crystal (blue) at 560nm. As a conductor, copper has a complex index of refraction, does not polarize incident light very strongly at Brewster’s angle and exhibits a gradual shift of retardance over the entire range of incident angles. For lead crystal, with its real-valued index of refraction of about 1.9, total polarization of incident light occurs at about 62° . Above this angle, no change in the phase relation of incident light occurs (both retardance components are at -90°), while below Brewster’s angle a phase difference of 180° is introduced.




Figure 8: Bispectral reflectivity measurements of pink fluorescent 3M Post-It® notes. The re-radiation matrix is shown for excitation wavelengths between 300nm and 780nm, and emission wavelengths from 380nm to 780nm, as 2D density plot and 3D graph. Data courtesy of Labsphere Inc.

Transfer from lower to higher energy levels primarily occurs in settings such as fluorescent lightsources, which are usually not modelled directly, but for which even highly realistic rendering systems just use the measured final combined emission spectrum. Common to both types of fluorescence is that they re-emit the incident light at different wavelengths within an extremely short time (typically 10^{-8} seconds).




Figure 9: Bispectral reflectivity measurements of pink fluorescent 3M Post-It® notes. In this 3D view the off-axis contribution had to be exaggerated in order to be properly visible, and both here and in figure 8 measurement noise is evident. Data courtesy of Labsphere Inc.

Instead of the reflectance spectra used for normal pigments, describing such a material requires knowledge of its re-radiation matrix, which encodes the energy transfer between different wavelengths. Such bispectral reflectance measurements are rather hard to come by; while “normal” spectrophotometers are becoming more and more common, the bispectral versions of such devices are by comparison very rare and in an experimental stage. Figures 8, 9 and 10 show three visualizations of a sample bispectral reflectance dataset. Manual design of such re-radiation matrices is much harder than explicit derivation of plain reflection spectra; while the latter is already not particularly easy, their effect is by comparison still quite predictable. Also, it is easy to maintain the energy balance of normal reflection spectra by ensuring that no component is greater than one; for a re-radiation matrix this translates to the more difficult condition that the integral over the area must not exceed one.

4. Previous work on Tone Reproduction

Reviews of tone reproduction operators have been carried out in previous years ^{36, 38}, and these also examine the HVS factors that influence the techniques.

Two types of tone reproduction operators can be used: *spatially uniform* (also known as *single-scale* or *global*) and *spatially varying* (also known as *multi-scale* or *local*). Spatially uniform operators apply the same transformation to every pixel regardless of their position in the image. A spatially




Figure 10: Bispectral reflectivity measurements of pink fluorescent 3M Post-It® notes. This graph shows **a**) the non-fluorescent reflection spectrum (the main diagonal of the re-radiation matrix in figure 9), **b**) the energy absorbed at higher wavelengths, **c**) the energy re-radiated at lower wavelengths and **d**) the resulting “reflection” spectrum. Note that the resulting spectrum is well over 1.0 in some areas. Data courtesy of Labsphere Inc.

uniform operator may depend upon the contents of the image as a whole, as long as the same transformation is applied to every pixel. Conversely, spatially varying operators apply a different scale to different parts of an image. A further aspect to tone reproduction is time. It should be noted that the above definitions do not account for temporal differences (such as adaptation over time), so we have included these under a separate category of *time dependent* tone reproduction operators. Figure 11 shows the development of tone reproduction methods and Table 4.3.1 gives an overview of tone reproduction methods published to date.

This section aims to provide an overview of the tone reproduction methods that have been published to date. Several of these methods are shown in greater detail. Tumblin and Rushmeier’s brightness preserving operator⁸³, which initially highlighted the importance of tone reproduction for computer graphics, is examined, as is Ward’s visibility preserving operator⁸⁵ — the basis for the development of Ferwerda et al.’s (and subsequently others’) time dependent method¹³.

4.1. Terminology

Throughout this paper all luminance measurements are given in cd/m^2 (candelas per square metre). The following terminology is also used:

- L Luminance
- w real-world
- d display

n a frame buffer value in the range $[0 \dots 1]$

4.2. Spatially uniform operators

Tumblin and Rushmeier Initial work on tone mapping in computer graphics was carried out by Tumblin and Rushmeier⁸³. They concentrated their task on preserving the viewer’s overall impression of brightness. Using a global operator they employed a psychophysical model of brightness perception developed by Stevens and Stevens⁷⁴ who produced a linear scale for brightness, where a 1 second exposure to a 5° white target of $1/(\pi \times 10^2) cd/m^2$ gives the unit of 1 *bril*. They showed that subjective brightness, B , grows as a power function of luminance

$$B = k(L - L_0)^\alpha$$

where k is a constant, L_0 is the minimum luminance that can be seen and α is an exponent between 0.333 and 0.49, depending on the level of adaptation. This relationship between target luminance and reported brightness is linear on a log-log scale.

This model of brightness perception is not valid for complex scenes but was chosen by Tumblin and Rushmeier due to its low computational costs. Their aim was to create a ‘hands-off’ method of tone reproduction in order to avoid subjective judgements. They created observer models — mathematical models of the HVS that include light-dependent visual effects while converting real-world luminance values to perceived brightness images. The real-world observer corresponds to someone immersed in the environment, and the display observer to someone viewing the display device. The hypothetical real-world observer would visually adapt to the luminance of the real-world scene, $L_{a(w)}$. The perceived brightness of this real-world luminance, L_w , can be calculated from

$$B_w = 10^{\beta(L_{a(w)})} (\pi \times 10^{-4} L_w)^\alpha(L_{a(w)})$$

where α and β are functions of the real-world adaptation level:

$$\alpha(l) = 0.4 \log_{10}(l) + 1.519$$

and

$$\beta(l) = -0.4(\log_{10}(l))^2 + 0.218 \log_{10}(l) + 6.1642$$

Tumblin and Rushmeier’s original luminance measurements were expressed in lamberts, but for the sake of consistency in this paper they have been converted to cd/m^2 (1 lambert = $\pi \times 10^{-4}$). (Note that in the above the term -0.218 rather than $+0.218$ has previously been used^{38, 16}, which appears to be a typographical error.)

Their tone reproduction operator converts the real-world luminances to the display values, which are chosen to match closely the brightness of the real-world image and the display image. If the image displayed on a CRT screen has a luminance value of L_d then its perceived brightness can be

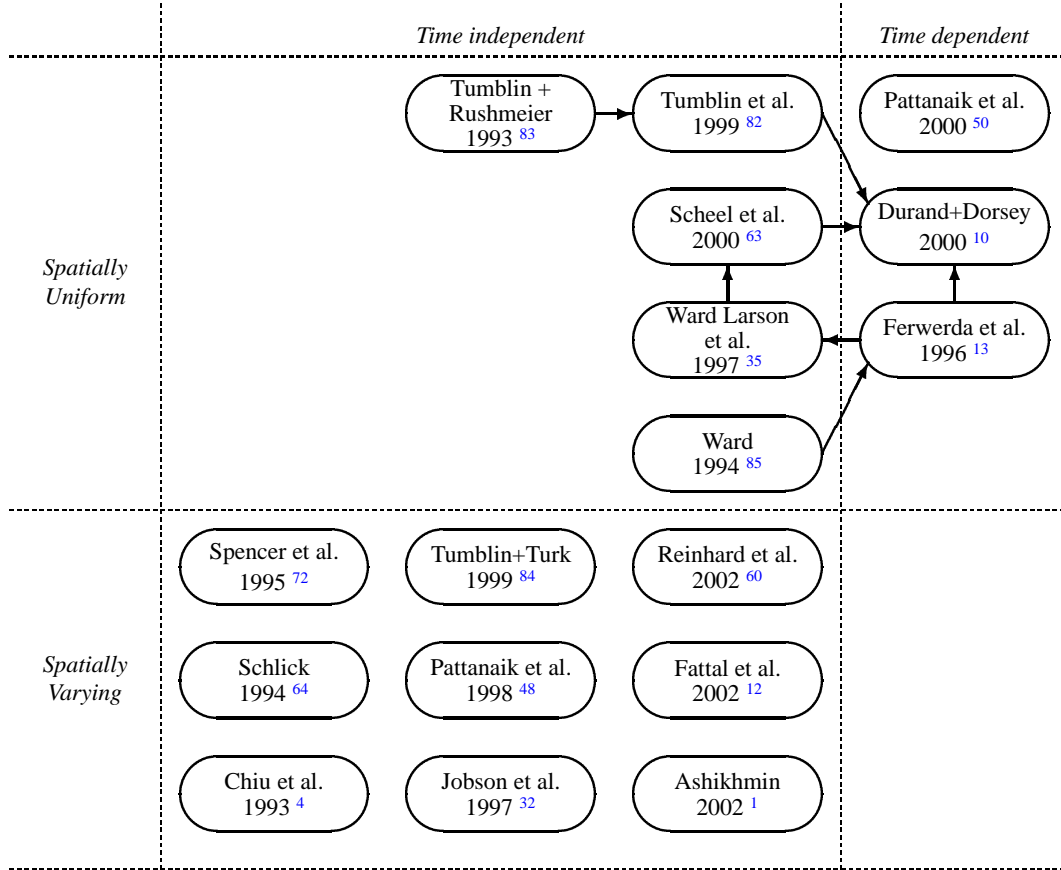


Figure 11: Taxonomy of Tone Reproduction Methods

expressed in a similar manner to the real-world luminance as:

$$B_d = 10^{\beta(L_{a(d)})} (\pi \times 10^{-4} L_d)^{\alpha(L_{a(d)})}$$

For the real-world luminance to match the display luminance, B_w must equal B_d . Using the aforementioned equations determining perceived brightnesses for real-world and display luminances, the display luminance required to match the real-world sensation can be obtained:

$$L_d = \frac{1}{\pi \times 10^{-4}} 10^{\frac{\beta_{a(w)} - \beta_{a(d)}}{\alpha_{a(d)}}} (\pi \times 10^{-4} L_w)^{\frac{\alpha_{a(w)}}{\alpha_{a(d)}}}$$

The the actual luminance produced for a frame buffer of n , with a gamma value between 2.2 and 2.5, is

$$L_d = L_{amb} + L_{dmax} n^\gamma$$

where L_{amb} is the ambient screen luminance and L_{dmax} is the maximum display value (approximately 100 cd/m^2). Inverting this will give us the required frame-buffer value to

produce the desired display luminance, ie.

$$n = \left[\frac{L_d - L_{amb}}{L_{dmax}} \right]^{1/\gamma}$$

giving the complete tone reproduction operator:

$$L_w = \left[\frac{10^{\frac{\beta_{a(w)} - \beta_{a(d)}}{\alpha_{a(d)}}} (\pi \times 10^{-4} L_w)^{\frac{\alpha_{a(w)}}{\alpha_{a(d)}}}}{\pi \times 10^{-4} L_{dmax}} - \frac{L_{amb}}{L_{dmax}} \right]^{1/\gamma}$$

If the display luminance falls outside the range of the frame-buffer $[0 \dots 1]$ then the frame-buffer value n will need to be clamped to fit this range.

This method is limited to greyscale and by the preservation of brightness at the expense of visibility in high dynamic scenes³⁵. It has been noted that as the operator can handle extremes of brightness, some images tend to appear too dark but this may work in its favour if the analysis of extreme lighting conditions is required³⁶.

Ward Ward's model ⁸⁵ dealt with the preservation of perceived contrast rather than brightness. Ward aimed to keep computational costs to a minimum by transforming real-world luminance values to display values through a scaling factor, concentrating on small alterations in luminance that are discernible to the eye. Based on a psychophysical contrast sensitivity model by Blackwell ² he took the view that the consequence of adaptation can be regarded as a shift in the absolute difference in luminance required for the viewer to notice the variation. Blackwell produced a comprehensive model of changes in visual performance due to adaptation level. Where the luminance difference, ΔL , is just noticeable between a target and a uniform background (when the observer has adapted to the luminance of the background, L_a) then:

$$\Delta L = 0.054(1.219 + L_a^{0.4})^{2.5}$$

Ward sought a scaling factor, m , between the display luminance, L_d , and the world luminance, L_w , so that

$$L_d = mL_w$$

To convert the real-world luminances to the display luminances so that the smallest discernible differences can be mapped then the assumption is made that

$$\Delta(L_{a(d)}) = m\Delta L(L_{a(w)})$$

where $\Delta(L_{a(d)})$ is the minimum discernible luminance change at $L_{a(d)}$, and as before, $L_{a(d)}$ is the display adaptation luminance and $L_{a(w)}$ is the real-world adaptation luminance. To find the scaling factor m then

$$m = \frac{\Delta L(L_{a(d)})}{\Delta L(L_{a(w)})} = \left[\frac{1.219 - L_{a(d)}^{0.4}}{1.219 - L_{a(w)}^{0.4}} \right]^{2.5}$$

is used.

This means that a Just Noticeable Difference (JND) in the real-world can be mapped as a JND on the display device. Ward assumes that the adaptation level is half the average radiance of the image, $L_{a(d)} = L_{dmax}/2$ as this is a close approximation for most applications. Thus, the final tone reproduction operator is:

$$L_w = \left[\frac{L_w}{L_{dmax}} \left[\frac{1.219 - L_{a(d)}^{0.4}}{1.219 - L_{a(w)}^{0.4}} \right]^{2.5} - \frac{L_{amb}}{L_{dmax}} \right]^{1/\gamma}$$

This approach is useful for displaying scenes where visibility analysis is crucial, such as emergency lighting, as it preserves the impression of contrast. It is also less computationally expensive than Tumblin and Rushmeier's operator but the use of a linear scaling factor causes very high and very low values to be clamped and correct visibility is not maintained throughout the image ³⁵. It should also be noted that Blackwell's experiments were conducted in near-perfect laboratory conditions and therefore do not take into consideration the complexities of typical workplace viewing conditions.

Ward Larson, Rushmeier and Piatko Further work by Ward Larson et al. ³⁵ presented a histogram equalisation technique for reproducing perceptually accurate tones in high dynamic display scenes, extending earlier work by Ward ⁸⁵ and Ferwerda et al. ¹³. They took object visibility and image contrast as their main focus, and considered a secondary goal of recreating the viewer's subjective response so that their impression of the real and virtual scenes were consistent ³⁵. Their model takes a histogram of scene brightnesses (the log of luminances averaged over 1° areas) which correspond with foveal adaptation levels for possible points in an image. A histogram and a cumulative distribution function are then obtained. Finally, Ferwerda et al.'s threshold sensitivity data is used to compress the original dynamic range to that of the display device, subject to the contrast sensitivity limits of the eye. Although this method is spatially uniform, spatial variation was introduced through the use of models for glare, acuity and chromatic sensitivity to increase perceptual fidelity.

Tumblin, Hodgkins and Guenter In 1999 Tumblin et al. ⁸² produced two new tone reproduction operators by imitating some of the HVS's visual adaptation processes. The first, a layering method, builds a display image from several layers of lighting and surface properties. This is done by dividing the scene into layers and compressing only the lighting layers while preserving the scene reflectances and transparencies, thus reducing contrast while preserving image detail. Their compression function follows the work of Schlick ⁶⁴. This method only works for synthetic images where layering information from the rendering process can be retained.

The second, a foveal method, interactively adjusts to preserve the fine details in the region around the viewer's gaze (which the viewer directs with a mouse) and compresses the remainder. In this instance their final tone reproduction operator is a revised version of the original Tumblin and Rushmeier operator, also building on the work of Ferwerda ¹³ and Ward ⁸⁵.

Both of these operators are straightforward in implementation and are not computationally expensive. The layering method is suited to static, synthetic scenes (displayed or printed) and the foveal method to interactive scenes (requiring a computer display).

Scheel, Stammering and Seidel Scheel et al. ⁶³ developed algorithms that permitted tone reproduction for interactive applications by representing luminances as a texture, allowing walkthroughs of large scenes where the tone reproduction can be adjusted frame-by-frame to the current view of the user, and focusing on tone reproduction for global illumination solutions obtained by radiosity methods. Due to interactivity, updates in tone mapping are required to account for changes in view point and viewing direction, and new factors need to be incorporated into the tone reproduction

operator, such as computational speed and adaptation determination. Tumblin et al.'s foveal method⁸² was interactive to an extent, but relied on pre-computed still images where the fixation point of the viewer could change, but an interactive walkthrough was not possible.

Spatially uniform operators were chosen due to computational efficiency, and Scheel et al. based their work on operators developed by Ward⁸⁵ and Ward Larson et al.³⁵. It uses a centre-weighted average to determine the probability of the user's focus. The adaptation levels are computed using samples obtained through ray-tracing, and the luminance of every vertex is held in texture co-ordinates. This can then be updated frame-by-frame. This method of tone reproduction provided a new level of interactivity, but did not take into consideration adaptation over time.

4.3. Spatially varying operators

Chiu, Herf, Shirley, Swamy, Wang and Zimmerman Chiu et al.'s⁴ investigation into global operators led them to believe that the solution should be local instead, as applying the same mapping to each pixel could produce incorrect results. They deliberately did not incorporate adaptation issues or psychophysical models into their operator; rather they experimented with a method of spatially varying image mapping. They showed that some pixels in an original image may have differing intensities in the display image dependent on their spatial position. As the HVS is more sensitive to relative as opposed to absolute changes in luminance they developed a spatially non-uniform scaling function for high contrast images. They based their work on the argument that the eye is more sensitive to reflectance than luminance, so that slow spatial variation in luminance may not be greatly perceptible. The implication is that images with a wider dynamic range than the display device can be displayed without much noticeable difference if the scaling function has a low magnitude gradient. By blurring the image to remove high frequencies, and inverting the result, the original details can be reproduced, but reverse intensity gradients appear when very bright and very dark areas are in close proximity³⁸.

Due to the fact that it is a local operator, this model is also computationally demanding. It is also a 'hands-on' approach, based purely on experimental results and therefore does not have the advantages of the more robust, theoretical basis of other tone reproduction operators.

Schlick Schlick⁶⁴ proposed several methods based on rational tone reproduction, but these were of an experimental nature only and also did not employ psychovisual models of the HVS, but concentrated on improving computational efficiency and simplifying parameters. He used a first degree rational polynomial function to map real-world luminances to display values, a function which worked satisfactorily when applied uniformly to all pixels in an image. His attempts

at accounting for local adaptation were less successful, but nonetheless worthy of mention in their development of Chiu et al.'s⁴ ideas.

Spencer, Shirley, Zimmerman and Greenberg Spencer et al.'s⁷² contribution was to develop a method of increasing dynamic range through the inclusion of glare effects. The idea of adding glare effects was previously recognised by Nakamae et al.⁴² (although their algorithm did not account for the visual masking effects of glare).

Spencer et al. produced psychophysically-based algorithms for adding glare to digital images, simulating the flare and bloom seen around very bright objects, and carried out a psychophysical test to demonstrate that these effects increased the apparent brightness of a light source in an image. While highly effective, glare simulation is computationally expensive.

Jobson, Rahman and Woodell Jobson et al.³² based their method on the retinex theory of colour vision, producing a multi-scale version that achieved simultaneous dynamic range compression, colour consistency and lightness rendition, and tested it extensively on (real-world) test scenes and over 100 images. The retinex is a computational model of lightness and colour perception of human vision which estimates scene reflectances, and Jobson et al. modified it to perform in a functionally similar manner to human visual perception. While this method worked well with their 24-bit RGB test images, they expressed the need for refinement of the method for images with greater maximum contrasts. Also, problems arose with scenes dominated by one colour as they violated the retinex "gray-world" assumption that the average reflectances are equal in the three spectral colour bands.

Pattanaik, Ferwerda, Fairchild and Greenberg This model developed by Pattanaik et al.⁴⁸ was based on a multi-scale representation of pattern, luminance, and colour processing in the HVS and addressed the problems of high dynamic range and perception of scenes at threshold and supra-threshold levels. They provide a computational model of adaptation and spatial vision for realistic tone reproduction. There are two main parts: the *visual model*, which processes an input image to encode the perceived contrasts for the chromatic and achromatic channels in their band-pass mechanism; and the *display model*, which takes the encoded information and outputs a reconstructed image. Although it is still computationally demanding, the model takes chromatic adaptation into account. However, this method is susceptible to strong halo effects⁸¹. Although it was designed as a solution towards the tone reproduction problems of wide absolute range and high dynamic range scenes, it is a general model that can be applied across a number of areas such as image quality metrics, image compression methods and perceptually-based image synthesis algorithms⁴⁸.

Tumblin and Turk Another method of tone reproduction was developed by Tumblin and Turk ⁸⁴ in 1999. This was the Low Curvature Image Simplifier (LCIS) method, a versatile method that could accept input from synthetic sources or real-world image maps, and either static or interactive. Similar to Tumblin et al.'s ⁸² layering and foveal approaches, the LCIS separates the input scene into large features and fine details, compressing the former and preserving the latter. The idea stems from art where an initial sketch outlines the main structure of a picture, with details and shadings filled in later. The LCIS uses a form of anisotropic diffusion to define the fine details by scene boundaries and smooth shading. This provides a high amount of subtle detail, avoids halo artifacts, and claims moderate computational efficiency ⁸¹.

Ashikhmin Ashikhmin has recently produced a new tone mapping operator which preserves image details and also provides enough information regarding absolute brightness in a low dynamic range image. It makes use of a basic functional model of human luminance perception, and is simple in implementation and moderate in computational expense. A locally-linear mapping (using a TVI function — see section 4.3.1) is computed for areas of uniform contrast, thus providing a similar level of detail preservation throughout the image.

Fattal, Lischinski and Wermann Recent work by Fattal et al. ¹² presents a new method for displaying HDR scenes. Their method is computationally efficient and conceptually simple, and is based on attenuating the magnitudes of the large luminance gradients that exist in HDR scenes, compressing large gradients and preserving fine details. The changes in intensity are identified and the larger gradients are reduced and a low dynamic range image is produced. They do not make an attempt at perceptual accuracy, but instead offer an effective, fast and easy-to-use form of tone reproduction.

Reinhard, Stark, Shirley and Ferwerda Reinhard et al. ⁶⁰ base their recent method on photographic practice, resulting in a technique designed to suit a wide variety of images. They use a photographic approach known as the Zone System, which divides a scene into 11 print zones ranging from pure black (zone 0) to white (zone X). A luminance reading is taken for a subjectively-defined middle-grey tone and this is typically mapped to zone V. Readings are taken for light and dark regions and a dynamic range can be determined, and an appropriate choice for middle-grey ensures that the maximum possible detail is retained.

Reinhard et al. attempt to manage tone reproduction choices rather than mimic the photographic process. A scaling similar to setting photographic exposure is applied, then highlighting and darkening selected regions (known as “dodging and burning” in photography) allows contrast to be controlled locally in the image over regions bounded by high

contrasts. They tested their method against existing tone reproduction operators with a broad range of HDR images. The model is confined to photographic dynamic range as opposed to the wider computer graphics measure, but it is high in computational efficiency.

4.3.1. Time dependent operators

Ferwerda, Pattanaik, Shirley and Greenberg Ferwerda et al. ¹³ developed a model based on the concept of matching JNDs for a variety of adaptation levels, which accounts for changes in colour appearance, visual acuity and temporal sensitivity while preserving global visibility. Whereas previous work on accurate tone reproduction concentrated on overcoming display limitations, their model addresses parameters in addition to dynamic range. It accounts for both rod and cone response (changes in colour appearance) and takes into consideration the aspect of adaptation over time. They exploit the detectability of changes in background luminance in order to remove those frequencies imperceptible when adapted to real-world illumination. Detection threshold experiments were used as the basis of the work. This is a way of measuring visual sensitivity in a psychophysical manner. In the experiments a viewer adapts to a dark screen. On each trial a disk of light is flashed in the centre of the screen for a few hundred milliseconds. The viewer states whether or not they have seen it. If they have not, the intensity is increased on the next trial; if they have, the intensity is decreased on the next trial. This provides a detection threshold. By plotting the detection threshold against the corresponding background luminance, a threshold-versus-intensity (TVI) function is produced for both the display and the viewer.

The implementation of this model is based on Ward's JND concept. A global scale factor of $L_d = mL_w$ is used. The thresholds for cones are

$$\Delta L_w = \begin{cases} 0.1905 & \text{if } L_{a(w)} \leq 2.512 \times 10^{-3} \\ 0.0556L_{a(w)} & \text{if } L_{a(w)} \geq 79.433 \\ 0.190510^{(0.249 \log_{10}(L_{a(w)}) + 0.65)^{2.7}} & \text{otherwise} \end{cases}$$

and the thresholds for rods are

$$\Delta L_w = \begin{cases} 1.3804 \times 10^{-3} & \text{if } L_{a(w)} \leq 1.148 \times 10^{-4} \\ 0.4027L_{a(w)} & \text{if } L_{a(w)} \geq 0.0363 \\ 1.3804 \times 10^{-3} 10^{(0.405 \log_{10}(L_{a(w)}) + 1.6)^{2.18}} & \text{otherwise} \end{cases}$$

If the level of adaptation in for the real-world viewer falls in the photopic range (ie. above 10 cd/m^2) then a photopic tone-reproduction operator is applied (making use of the cone data), and if it falls in the scotopic range (ie. below 0.01 cd/m^2) then a scotopic tone reproduction operator is applied (making use of the rod data).

The required display luminance for the display device can

then be calculated by combining the photopic and scotopic display luminances using a parametric constant k

$$L_d = L_{d(p)} + k(L_{a(w)})L_{d(s)}$$

where k is in the range $[0..1]$ as the real-world adaptation level goes from the bottom and top of the mesopic range, ie. k is used to simulate the loss of rod sensitivity in the mesopic range. The achromaticity of rod-dominated vision results in the rod TVI function producing a grey-scale display luminance. When this is combined with the chromatic values produced by the cone TVI function, the loss of colour which occurs when the average luminance in a scene diminishes is displayed.

To reproduce the loss in visual acuity, Ferwerda et al. used data from psychophysical experiments that related the detectability of square wave gratings of different spatial frequencies to changes in background luminance. Using this data it is possible to determine what spatial frequencies are visible, and thereby eradicate any extraneous data in the image. To avoid ringing in the displayed image, a Gaussian convolution filter is used, thus removing frequencies which are not discernible to the real-world viewer.

Light and dark adaptation were also considered by adding a parametric constant b to the display luminance, the value of which changes over time

$$L_d = L_w \frac{\Delta L_d}{\Delta L_w} + b$$

No quantitative data was available to set b , so b was set such that

$$L_d(L_{a(w)}) = \text{constant over time}$$

which means that the overall luminance of the display will not change during the light adaptation process. The same procedure can be applied without change for dark adaptation.

This model is of particular importance due to the psychophysical model of adaptation that it adopts, and will prove useful for immersive display systems that cover the entire visual field so that the viewer's visual state is determined by the whole display ³⁸.

Pattanaik, Tumblin, Yee and Greenberg Pattanaik et al. ⁵⁰ offered a new time-dependent tone reproduction operator to rapidly create readily displayable colour image sequences from static, dynamic, real or synthetic input scenes. It follows the perceptual models framework proposed by Tumblin and Rushmeier with the addition of an adaptation model and appearance model to express retinal response and lightness and colour. They claimed that their work was novel in three ways - it is general, accepting time sequences of arbitrary scene intensities; it captures the appearance of widely varying amounts of adaptation; and it is firmly grounded in published research results from psychophysics, physiology and colour science ⁵⁰. The adaptation model computes retina-like

response signals (for rod and cone luminance and colour information) for each pixel in the scene. Using Hunt's static model of colour vision, they add time-dependent adaptation components to describe neural effects, pigment bleaching, regeneration and saturation effects. Their visual appearance model assumes that the real-world viewer will determine a 'reference white' and a 'reference black' and judge the appearance of any visual response against these standards. Assembling these models reproduces the appearance of scenes that evoke changes to visual adaptation. This operator is suitable for use in real-time applications as due to its global model of adaptation it does not require extensive processing. However, local adaptation - an important part of visual appearance - is not supported.

Durand and Dorsey Durand and Dorsey ¹⁰ presented an interactive tone mapping model which made use of visual adaptation knowledge. They also proposed extensions to Ferwerda et al.'s tone mapping operator and incorporated it into a model for the display of global illumination solutions and interactive walkthroughs. This model involves time-dependent tone mapping and light adaptation, and extends Ferwerda et al.'s by including a blue-shift for viewing night scenes and by adding chromatic adaptation. For the interactive implementation, work by Tumblin et al. ⁸² and Scheel et al. ⁶³ was used to take advantage of the observer's gaze, allowing a weighted average to be used. Photographic exposure metering used in photography is employed to better calculate the adaptation level. Loss of visual acuity is simulated in the same manner of Ferwerda et al. by use of a 2D Gaussian blur filter. The scene is rendered as normal, with interactivity introduced by tone mapping computed on the fly, accelerated by caching the function in Look-Up-Tables.

5. Previous Work in Spectral Rendering

In this section we first examine research efforts which aimed at implementing some of the special effects mentioned in section 3.2 which are only possible using a spectral rendering system. In subsection 5.4 we then discuss spectral rendering systems which have been presented in literature so far.

5.1. Raytracing with Polarization Parameters

There are three publications in computer graphics literature about this topic; one by Wolff and Kurlander ⁹¹, who demonstrated the concept for the first time, one by Tannenbaum et al. ⁷⁹, who concentrated on the rendering of anisotropic crystals and extended the techniques used by Wolff et al. , and one by Wilkie et al. ⁹⁰, who proposed using the Stokes vector formalism and incorporated fluorescence effects at the same time.

The main goal of these efforts was that of finding an appropriate way to describe, and perform calculations with,

Algorithm	Spatially		Time dependent	Strengths and Weaknesses
	Uniform	Varying		
Tumblin and Rushmeier 1993 ⁸³	✓			Preserves brightness. Does not preserve visibility or account for adaptation. Grey scale only.
Chiu et al 1993 ⁴		✓		Preserves local contrast. Ad hoc. Computationally demanding. Halo artifacts.
Ward 1994 ⁸⁵	✓			Preserves contrast - crucial for predictive lighting analysis. Clipping of very high and very low values. Does not consider complexities of typical workplace viewing.
Schlick 1994 ⁶⁴		✓		Speed and simplification of uniform and varying operators. Ad hoc - no perceptual accuracy.
Spencer et al. 1995 ⁷²		✓		Greater dynamic range through glare effects.
Ferwerda et al. 1996 ¹³	✓		✓	Accounts for changes in threshold visibility, colour appearance, visual acuity and sensitivity over time. Psychophysical model. Useful for immersive displays.
Ward Larson et al. 1997 ³⁵	✓			Histogram equalisation. Preserves local contrast visibility. Uses models for glare, colour sensitivity and visual acuity to increase perceptual realism.
Jobson et al. 1997 ³²		✓		Multi-scale retinex model - perceptually valid. Problems with monochrome scenes and maximum contrasts outside 24-bit RGB range.
Pattanaik et al. 1998 ⁴⁸		✓		Multi-scale psychophysical representation of pattern, luminance and colour processing resulting in increased perceptual fidelity. Does not incorporate temporal aspects.
Tumblin et al. 1999 ⁸²	✓			Layering method for static, synthetic images. Foveal method for interactive scenes.
Tumblin and Turk 1999 ⁸⁴		✓		LCIS method preserves subtle detail and avoids halo artifacts.
Pattanaik et al. 2000 ⁵⁰	✓		✓	Psychophysical operator with time-dependent adaptation and appearance models. Does not support local adaptation.
Scheel et al. 2000 ⁶³	✓			Interactive method using texturing hardware. Does not account for adaptation over time.
Durand and Dorsey 2000 ¹⁰	✓		✓	Interactive method with time-dependent adaptation and simulation of visual acuity and chromatic adaptation.
Ashikhmin 2002 ¹		✓		Preserves image details and absolute brightness information. Uses simple functional perceptual model. Simple to implement with moderate computational efficiency.
Fattal et al. 2002 ¹²		✓		Computationally efficient and simple operator. Does not attempt psychophysical accuracy.
Reinhard et al. 2002 ⁶⁰		✓		Method based on photographic technique. Suits a wide variety of images, but limited to photographic dynamic range rather than the wider computer graphics dynamic range.

Table 1: Tone Reproduction Operators: summary of attributes

polarized light; both earlier groups of authors settled for a notation suggested by standard reference texts from physics literature.

5.1.1. Coherency Matrices

The formalism to describe polarized light used by both Wolff and Tannenbaum is that of *coherency matrices* (CM for short); this technique was introduced by Born and Wolf³. Similar to the discussion in section 3.5, one considers a monochromatic wave propagating along the positive Z -axis, and treats the E -field vector as consisting of two orthogonal components in X and Y directions. The phase relationships between the two can exhibit anything between full correlation (fully coherent or polarized light) to totally uncorrelated behaviour (incoherent or unpolarized light). As derived in detail by Tannenbaum et al.⁷⁹, the coherency state of such a monochromatic wave can be expressed in a matrix J of the form

$$J = \begin{pmatrix} J_{xx} & J_{xy} \\ J_{yx} & J_{yy} \end{pmatrix} = \begin{pmatrix} \langle E_x E_x^* \rangle & \langle E_x E_y^* \rangle \\ \langle E_y E_x^* \rangle & \langle E_y E_y^* \rangle \end{pmatrix}$$

where E_x and E_y are the time average of the E -field vectors for the X and Y directions, respectively, and E_x^* denotes the complex conjugate of E_x . The main diagonal elements – J_{xx} and J_{yy} – are real valued, and the trace $T(J) = J_{xx} + J_{yy}$ of the matrix represents the total light radiation of the wave, while the complex conjugates J_{xy} and J_{yx} represent the correlation of the X and Y components of E . For fully polarized light, these components are fully correlated and $|J|$ vanishes.

It is worth noting that coherency matrices – while they might seem rather different at first glance – encode the same information as equation (5). Born and Wolf³ also discuss how one can compute the elements of a coherency matrix from actual irradiance measurements, which is an important point if one wants to eventually validate computer graphics models that use polarization in their calculations.

5.1.2. Coherency Matrix Modifiers

Besides being able to describe a ray of light, it is also necessary to process the interaction of light with a medium, as outlined in section 3.4.1. Such filtering operations on polarized light described by a coherency matrix can be performed by using *coherency matrix modifiers*, or CMM. Tannenbaum et al. brought this approach, which was originally introduced by Parrent et al.⁴⁷, to the computer graphics world for use in their rendering system.

CMMs have the form of a complex-valued 2×2 matrix. If all participating elements have the same reference coordinate system, such matrices can be applied to a given coherency matrix J in the sequence of $J_p = \mathcal{M}_p J \mathcal{M}_p^\dagger$, where \mathcal{M}_p^\dagger signifies the conjugate transpose of \mathcal{M}_p . If the modifier and the coherency matrix are not in the same coordinate system, an appropriate transformation – as discussed in section 2.3 of Tannenbaum et al.⁷⁹ – has to be applied first.

5.1.3. The Stokes Vector Formalism

An alternative description⁹⁰ for polarized radiation with somewhat simpler mathematical characteristics is that of *Stokes parameters*. This description, while equivalent to coherency matrices, has the property of using only real-valued terms to describe all polarization states of optical radiation, and has an – also noncomplex – corresponding description of ray weights in the form of Müller matrices⁶⁸.

It has to be kept in mind that – similar to coherency matrices – both Stokes parameters and Müller matrices are meaningful only when considered within their own local reference frame; the main effect of this is that in a rendering system not only light, but also filters are *oriented* and have to store an appropriate reference in some way. However, for the sake of brevity this spatial dependency is omitted in our following discussion except in the section about matrix realignment.

Three real-valued parameters are required to describe a general polarization ellipse, but the *Stokes vector* notation defined by

$$\begin{aligned} E_{n,0} &= \kappa(V_x^2 + V_y^2) & [W \cdot m^{-2}] \\ E_{n,1} &= \kappa(V_x^2 - V_y^2) \\ E_{n,2} &= \kappa(2V_x^2 \cdot V_y^2 \cdot \cos\gamma) \\ E_{n,3} &= \kappa(2V_x^2 \cdot V_y^2 \cdot \sin\gamma) \end{aligned} \quad (7)$$

has proven itself in the optical measurements community, and has the key advantage that the first component of this 4-vector is the unpolarized intensity of the light wave in question (i.e. the same quantity that a nonpolarizing renderer uses). Components 2 and 3 describe the preference of the wave towards linear polarization at zero and 45 degrees, respectively, while the fourth encodes preference for right-circular polarization. While the first component is obviously always positive, the values for the three latter parameters are bounded by $[-E_{n,0}, E_{n,0}]$; e.g. for an intensity $E_{n,0} = 2$, a value of $E_{n,3} = -2$ would indicate light which is totally left circularly polarized.

The – at least in comparison to coherency matrices – much more comprehensible relationship between the elements of a Stokes vector and the state of the wavetrain it describes can prove beneficial during the debugging stage of a polarizing renderer, since it is much easier to construct verifiable test cases.

5.1.4. Müller Matrices

Müller matrices (MM for short) are the data structure used to describe a filtering operation by materials that are capable of altering the polarization state of incident light represented by a Stokes vector. The general modifier for a 4-vector is a 4×4 -matrix, and the structure of the Stokes vectors implies that the elements of such a matrix correspond to certain physical filter properties. As with Stokes vectors, the better

comprehensibility of these real-valued data structures is of considerable benefit during filter specification and testing.

The degenerate case of MM is that of a nonpolarizing filter; this could equally well be described by a simple reflection spectrum. For such a filter the corresponding MM is the identity matrix. A more practical example is the MM of an ideal linear polarizer T_{lin} , where the polarization axis is tilted by an angle of ϕ against the reference coordinate system of the optical path under consideration, and which has the form of

$$T_{lin}(\phi) = \frac{1}{2} \begin{bmatrix} 1 & \cos 2\phi & \sin 2\phi & 0 \\ \cos 2\phi & \cos^2 2\phi & \sin 2\phi \cdot \cos 2\phi & 0 \\ \sin 2\phi & \sin 2\phi \cdot \cos 2\phi & \sin^2 2\phi & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

For the purposes of physically correct rendering it is important to know the MM which is caused by evaluation of the Fresnel terms in equation (6). For a given wavelength and intersection geometry (i.e. specified index of refraction and angle of incidence), the resulting terms F_{\perp} , F_{\parallel} , δ_{\perp} and δ_{\parallel} have to be used as

$$T_{Fresnel} = \begin{bmatrix} A & B & 0 & 0 \\ B & A & 0 & 0 \\ 0 & 0 & C & -S \\ 0 & 0 & S & C \end{bmatrix},$$

where $A = (F_{\perp} + F_{\parallel})/2$, $B = (F_{\perp} - F_{\parallel})/2$, $C = \cos(\delta_{\perp} - \delta_{\parallel})$ and $S = \sin(\delta_{\perp} - \delta_{\parallel})$; $\delta_{\perp} - \delta_{\parallel}$ is the total retardance the incident wavetrain is subjected to.

5.1.5. Filter Rotation

In order to correctly concatenate a filter chain described in section 3.4.1 (which basically amounts to matrix multiplications of the MMs in the chain), we have to be able to re-align a MM to a new reference system, which amounts to rotating it along the direction of propagation to match the other operands.

Contrary to first intuition, directional realignment operations are *not* necessary along the path of a concatenated filter chain; the retardance component of a surface interaction is responsible for the alterations that result from changes in wavetrain direction.

Since in the case of polarized light a rotation by an angle of ϕ can only affect the second and third components of a Stokes vector (i.e. those components that describe the linear component of the polarization state), the appropriate rotation matrix $M(\phi)$ has the form of

$$M(\phi) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos 2\phi & \sin 2\phi & 0 \\ 0 & -\sin 2\phi & \cos 2\phi & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (8)$$

Matrices of the same form are also used to rotate Müller matrices. In order to obtain the rotated version of a MM $T(0)$,

$M(\phi)$ has to be applied in a way similar to that shown for CMMs, namely $T(\phi) = M(-\phi) \cdot T(0) \cdot M(\phi)$.

Apart from being useful to re-align a filter through rotation, the matrix M given in equation (8) is also the MM of an ideal *circular retarder*. Linearly polarized light entering a material of this type will emerge with its plane of polarization rotated by an angle of ϕ ; certain materials, such as crystal quartz or dextrose, exhibit this property, which is also referred to as *optical activity*.

5.2. Reflection Models which take Polarization into Account

Apart from the case of perfect specular reflection and refraction, which is already adequately covered by the Fresnel terms discussed in section 3.4.1 and the modified Cook-Torrance model used by Wolff et al., only one other surface model proposed so far, namely that of He et al.²⁵, attempts to consider polarization effects. As could be inferred from the results section of their paper, the high complexity of their surface model apparently led the authors to only implement a simpler, non-polarizing version in practice, and to contain themselves with just providing the theoretical derivation of the polarization-aware model in the text.

5.3. Rendering of Fluorescence Effects

Glassner was the first graphics researcher who investigated the rendering of fluorescence phenomena¹⁷. The main focus of his work was centered around the proper formulation of the rendering equation in the presence of phosphorescence and fluorescence, and he provided striking results generated with a modified version of the public domain raytracer *rayshade*. Sadly, this work was not followed up, nor was the modified version of rayshade made public. Also, we are not aware of any work that aims at considering the inclusion of fluorescence effects in sophisticated reflectance models.

Wilkie et al.⁹⁰ presented a hybrid renderer which is both capable of including polarization and fluorescence effects at the same time.

Since fluorescence is a material property, its description only affects the filter data structure. Specifically, for a system which uses n samples to represent spectra, filter values of fluorescent substances have to be a re-radiation matrix (RRM) of $n \times n$ elements.

The fact that those types of fluorescence effects which a rendering system might want to model only cause light to be re-radiated at *lower* wavelengths (see section 3.6) allows one to only consider the lower half of this matrix. As pointed out by Glassner¹⁷ this in turn has the very beneficial consequence that one does not have to perform equilibrium calculations at all wavelengths for each light-surface interaction, since an ordinary matrix multiplication suffices in this case.

5.4. Spectral Rendering Systems

Somewhat surprisingly, the usage of spectral methods in the rendering research community has been comparatively infrequent and limited to certain research projects, although this is improving slowly.

None of the current standard general-purpose rendering toolkits which are publicly available in source code – such as Radiance⁸⁷ or RenderPark – support it, but there is hope that future projects might provide spectral rendering capabilities.

5.4.1. The Cornell Image Synthesis Testbed

Probably the earliest full spectral renderer was the system developed by Hall and Greenberg²⁴. It was able to perform colour computations at up to 1nm intervals through the visual range, and could be programmed to use any number of arbitrarily spaced samples as well as RGB and CIE XYZ colour models. The best efficiency vs. speed trade-off was apparently achieved at 9 samples, which incurred an execution time penalty of just two percent over the RGB version of the renderer, with greatly improved colour fidelity in the output.

It is worth noting that the Cornell image synthesis testbed was also revolutionary in another way that pertains to this STAR report: it stored its result images as (comparatively) high dynamic range CIE XYZ colour space images, something that it took other systems a long time to emulate.

5.4.2. Wavefront Tracking

In the course of his PhD work, Collins^{6,5} developed a spectral rendering system which he used as a basis for the development of a precursor to modern photon tracing global illumination techniques called *wavefront tracking*. His system was capable of rendering dispersion effects, and used the full Fresnel terms for spectral evaluation of surface reflectancies. The number of spectral samples taken was an option left to the user of the system, and the text remains inconclusive if the author found a reliable upper margin for this value.

Collins also used his spectral renderer to explore the visual properties of compound eyes⁷. Since insects have different colour response functions and can see in a different range of the electromagnetic spectrum, one needs spectral images to perform such comparisons to human vision.

5.4.3. Linear Color Representations for Full Spectral Rendering

Peercy⁵¹ was amongst the first to suggest an alternative to the direct spectral sampling technique used by researchers up to then. He argued that brute force sampling is less than efficient in most cases, and proposed the use of colour representations based on sets of carefully chosen basis functions. Using this technique he was able to satisfactorily describe a spectral rendering process by using as little as four or five spectral samples.

He later attempted to utilize some of these ideas for an interactive spectral rendering system which was developed up to prototype level by SGI⁵²; further development of this system was unfortunately discontinued due to lack of customer interest.

5.4.4. Adaptive Spectral Rendering with Perceptual Control

Iehl et al.³⁰ have presented a hybrid stochastic rendering system which was designed from the ground up to support spectral rendering. The system features adaptive representation of spectral intensities and attempts to operate on the least costly spectral representation it can use to meet given accuracy constraints.

5.4.5. Hybrid Spectral Rendering

Sun et al.^{78, 76, 75} proposed a novel approach to spectral rendering that aims to overcome both the limitations of the fixed sampling technique and the overhead associated with basis functions.

Low-order basis functions are used for representation of the overall spectral shape, and spectral spikes are maintained separately. In this way a very high accuracy is maintained even though just a few coefficients have to be stored. With his system, the author demonstrated the computation of realistic reflection models and dispersion effects at comparatively very low cost.

5.4.6. The Utah Spectral Renderer

The graphics group originally centered around the University of Utah has used a spectral rendering system for a number of high-profile papers about photorealistic rendering of natural phenomena such as skylight⁵⁶, water⁵⁷ and night skies³¹, and in the case of skylight rendering has even made parts of the source available on the web.

Although their entire rendering system has not been released, the research conducted by this group is a breakthrough for spectral rendering insofar as they are the first to routinely quote spectral radiances and reflectivities in the appendices of their papers.

Their publications do not go into details which type of spectral representation is being used internally by their system, but from the published source code fragments it can be inferred that they at least partially use a fixed sampling approach.

5.4.7. Others

Raso et al.⁵⁹ proposed to use piecewise cubic polynomials to approximate spectral distributions, thereby reducing inter-reflection computations to polynomial multiplications.

Icart et al.^{27, 28, 29} researched physical models for the interaction of light with matter, and came up with very sophisticated results and derivations of reflectancy models. The

rendering system developed for their experiments used the four fixed wavelength approach proposed by Meyer ³⁹.

Evans et al. ¹¹ proposed a wavelength dependent Monte Carlo rendering technique. The stratified wavelength clustering (SWC) strategy carries several wavelength stratified radiance samples along each light transport path. The cluster is split into several paths or degraded into a single path only if a specular refraction at the surface of a dispersive material is encountered along the path. The overall efficiency of this strategy is high since the fraction of clusters that need to be split or degraded in a typical scene is low, and also because specular dispersion tends to decrease the source colour variance, offsetting the increased amortized cost of generating each path.

Rougeron et al. ⁶² proposed an adaptive representation scheme for spectral data that minimises the colorimetric error incurred through rendering operations. The spectra are projected to a set of hierarchical basis functions, which leads to a representation through binary trees. Based on this representation, an adaptive algorithm which controls the absolute error of refinement and merge steps during rendering was presented. In later work ³⁰, one of the authors proceeded to incorporate perceptual measurements into this approach.

Gondek et al. ²⁰ presented a wavelength based bidirectional reflectance function developed for use in realistic image synthesis. They employed a geodesic sphere to represent the BRDF, and implemented a virtual goniospectrophotometer by using a Monte Carlo ray tracer to cast rays into a surface. They also used an optics model that incorporates phase in the ray tracer to simulate interference effects. An adaptive subdivision technique was applied to elaborate the data structure from rays scattered into the hemisphere above the surface. The wavelength based BRDF and virtual goniospectrophotometer was then utilized to analyze and make pictures of thin films, idealized pigmented materials, and pearlescent paints.

6. Conclusion

There is still a great deal of work to be undertaken in the area of realistic image synthesis. While most efforts seem to be directed towards improvements of rendering algorithms, the two key areas of spectral rendering and tone reproduction both still have important questions open for future research.

For spectral rendering these issues are the problem which spectral representation is optimal for a given purpose, and how to best incorporate advanced properties of light such as polarization, phosphorescence and fluorescence into a rendering system without huge losses in efficiency.

The second key area for improvements is that of tone reproduction. Until advances in hardware provide us with a more advanced form of display we will have to depend on tone reproduction operators to deliver the desired perceptual

effect. Evolution of display technology has provided us with flat screen LCD displays and micro-mirror projection systems, but these have still to become commonplace, and have disadvantages of their own. In the case of LCD displays, limitations are imposed due to shortcomings involving angular dependence, temperature, channel constancy, resolution and a lack of control of gamma. The development of a High Dynamic Range viewer will allow for testing of existing tone reproduction operators, allowing us to apply the most effective models for our purpose. (Greg Ward has produced an experimental high contrast stereoscopic display with a contrast ratio of 5000:1.) Although our display capabilities are limited, it is important to ensure that the information is stored in a relevant device-independent representation so that none of the HDR information is lost, thus preserving display options. Formats such as the SGI LogLuv TIFF, which can hold 38 orders of magnitude in its 32-bit version, have been recommended ^{34, 86}.

The present path of adopting psychophysical methods would seem to be the best way forward, but knowledge of the HVS is still limited, and modelling its characteristics for the purpose of tone mapping will be complex and time consuming. The lack of comprehensive image metrics in graphics also limits the study. At present, the answer is to use the most appropriate method for the situation. Depending on requirements, a number of different operators are available for use and they must be selected on the premise of the 'best tool for the job'. There is undoubtedly a need for the validation of tone reproduction operators, preferably through psychophysical comparison.

Acknowledgements

Thanks to Greg Ward for a helpful discussion on tone reproduction, Robert Tobler for general input on the topic of spectral rendering, and Katharina Horrak and Peter Gerhardus for clarifying discussions on the topic of fluorescence.

As usual, the Vienna group would also like to thank The Master for His Spiritual Guidance during the preparation phase of this STAR report.

References

1. M. Ashikhmin. A tone mapping algorithm for high contrast images. In *13th Eurographics Workshop on Rendering*. Eurographics, June 2002. 10, 15
2. H. R. Blackwell. *An Analytical Model for Describing the Influence of Lighting Parameters upon Visual Performance*, volume 1: Technical Foundations. Comission Internationale De L'Eclairage, 1981. 11
3. Max Born and Emil Wolf. *Principles of Optics*. The Macmillan Company, 1964. 5, 16
4. K. Chiu, M. Herf, P. Shirley, S. Swamy, C. Wang, and

- K. Zimmerman. Spatially nonuniform scaling functions for high contrast images. In *Graphics Interface '93*, pages 245–253, Toronto, Ontario, Canada, May 1993. Canadian Information Processing Society. 10, 12, 15
5. Steven Collins. Rendering crystal glass. Technical Report TCD-CS-94-20, Department of Computer Science, University of Dublin, Trinity College, 31 December 1994. Fri, 18 Apr 1997 12:38:03 GMT. 5, 18
 6. Steven Collins. *Wavefront Tracking for Global Illumination Solutions*. PhD thesis, Trinity College, Dublin, Ireland, November 1996. Available from <http://isg.cs.tcd.ie/scollins/work.html>. 18
 7. Steven Collins. Reconstructing the visual field of compound eyes. In Julie Dorsey and Philipp Slusallek, editors, *Eurographics Rendering Workshop 1997*, pages 81–92, New York City, NY, June 1997. Eurographics, Springer Wien. ISBN 3-211-83001-4. 18
 8. R. L. Cook and K. E. Torrance. A reflectance model for computer graphics. *Computer graphics*, Aug 1981, 15(3):307–316, 1981. 6
 9. P. M. Deville, S. Merzouk, D. Cazier, and J. C. Paul. Spectral data modeling for a lighting application. *Computer Graphics Forum*, 13(3):C/97–C/106, 1994. 4
 10. Fredo Durand and Julie Dorsey. Interactive tone mapping. In *Rendering Techniques 2000: 11th Eurographics Workshop on Rendering*, pages 219–230. Eurographics, June 2000. ISBN 3-211-83535-0. 10, 14, 15
 11. Glenn F. Evans and Michael D. McCool. Stratified wavelength clusters for efficient spectral monte carlo rendering. In *Graphics Interface*, pages 42–49, June 1999. 19
 12. Raanan Fattal, Dani Lischinski, and Micheal Werman. Gradient domain high dynamic range compression. In *Proceedings of ACM SIGGRAPH 2002*, Computer Graphics Proceedings, Annual Conference Series. ACM Press / ACM SIGGRAPH, July 2002. 10, 13, 15
 13. James A. Ferwerda, Sumant Pattanaik, Peter S. Shirley, and Donald P. Greenberg. A model of visual adaptation for realistic image synthesis. In *Proceedings of SIGGRAPH 96*, Computer Graphics Proceedings, Annual Conference Series, pages 249–258, New Orleans, Louisiana, August 1996. ACM SIGGRAPH / Addison Wesley. ISBN 0-201-94800-1. 2, 9, 10, 11, 13, 15
 14. James A. Ferwerda, Sumanta N. Pattanaik, Peter S. Shirley, and Donald P. Greenberg. A model of visual masking for computer graphics. In *Proceedings of SIGGRAPH 97*, Computer Graphics Proceedings, Annual Conference Series, pages 143–152, Los Angeles, California, August 1997. ACM SIGGRAPH / Addison Wesley. ISBN 0-89791-896-7.
 15. R. Geist, O. Heim, and S. Junkins. Color representation in virtual environments. *Color Research and Application*, 21(2):121–128, April 1996. 4
 16. S. Gibson. *Efficient Radiosity Simulation using Perceptual Metrics and Parallel Processing*. Phd thesis, University of Manchester, September 1998. 9
 17. Andrew Glassner. A model for fluorescence and phosphorescence. In *Fifth Eurographics Workshop on Rendering*, pages 57–68, Darmstadt, Germany, June 1994. Eurographics. 17
 18. Andrew S. Glassner. How to derive a spectrum from an RGB triplet. *IEEE Computer Graphics and Applications*, 9(4):95–99, July 1989. 4
 19. Andrew S. Glassner. *Principles of Digital Image Synthesis*. Morgan Kaufmann, San Francisco, CA, 1995. 3, 4, 5, 7
 20. Jay S. Gondek, Gary W. Meyer, and Jonathan G. Newman. Wavelength dependent reflectance functions. *Computer Graphics*, 28(Annual Conference Series):213–220, July 1994. 19
 21. Donald P. Greenberg, Kenneth E. Torrance, Peter S. Shirley, James R. Arvo, James A. Ferwerda, Sumanta Pattanaik, Eric P. F. Lafortune, Bruce Walter, Sing-Choong Foo, and Ben Trumbore. A framework for realistic image synthesis. In *Proceedings of SIGGRAPH 97*, Computer Graphics Proceedings, Annual Conference Series, pages 477–494, Los Angeles, California, August 1997. ACM SIGGRAPH / Addison Wesley. ISBN 0-89791-896-7.
 22. Roy Hall. *Illumination and Color in Computer Generated Imagery*. Springer-Verlag, New York, NY, 1989. 3, 4
 23. Roy Hall. Comparing spectral color computation methods. *IEEE Computer Graphics and Applications*, 19(4), July 1999. 3
 24. Roy A. Hall and Donald P. Greenberg. A testbed for realistic image synthesis. *IEEE Computer Graphics and Applications*, 3(8):10–20, November 1983. 3, 18
 25. Xiao D. He, Kenneth E. Torrance, François X. Sillion, and Donald P. Greenberg. A comprehensive physical model for light reflection. *Computer Graphics*, 25(4):175–186, July 1991. 17
 26. R. W. G. Hunt. *The Reproduction of Colour in Photography, Printing and Television*. Fountain Press, Tolworth, 5th edition, 1995. 2
 27. Isabelle Icart and Didier Arquès. An approach to geometrical and optical simulation of soap froth. *Computers and Graphics*, 23(3):405–418, June 1999. 18
 28. Isabelle Icart and Didier Arquès. An illumination model

- for a system of isotropic substrate - isotropic thin film with identical rough boundaries. In Dani Lischinski and Greg Ward Larson, editors, *Rendering Techniques '99*, Eurographics, pages 261–272. Springer-Verlag Wien New York, 1999. 18
29. Isabelle Icart and Didier Arquès. A Physically-Based BRDF model for multilayer systems with uncorrelated rough boundaries. In Bernard Péroche and Holly Rushmeier, editors, *Rendering Techniques '00*, Eurographics, pages 353–364. Springer-Verlag Wien New York, 2000. 18
 30. Jean Claude Iehl and Bernard Péroche. An adaptive spectral rendering with a perceptual control. *Computer Graphics Forum*, 19(3):291–300, August 2000. ISSN 1067-7055. 18, 19
 31. Henrik Wann Jensen, Frédo Durand, Michael M. Stark, Simon Premoze, Julie Dorsey, and Peter Shirley. A physically-based night sky model. In Eugene Fiume, editor, *SIGGRAPH 2001, Computer Graphics Proceedings*, Annual Conference Series, pages 399–408. ACM Press / ACM SIGGRAPH, 2001. 18
 32. D. J. Jobson, Z. Rahman, and G. A. Woodell. A multi-scale retinex for bridging the gap between color images and the human observation of scenes. *IEEE Transactions on Image Processing*, 6(7):965–976, July 1997. 10, 12, 15
 33. G. P. Können. *Polarized Light in Nature*. Cambridge University Press, 1985. 6
 34. Gregory Ward Larson. Logluv encoding for full-gamut, high-dynamic range images. *Journal of Graphics Tools*, 3(1):15–31, 1998. ISSN 1086-7651. 19
 35. Gregory Ward Larson, Holly Rushmeier, and Christine Piatko. A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Transactions on Visualization and Computer Graphics*, 3(4):291–306, October - December 1997. ISSN 1077-2626. 10, 11, 12, 15
 36. K. Matkovic, L. Neumann, and W. Purgathofer. A survey of tone mapping techniques. In *13th Spring Conference on Computer Graphics*, pages 163–170, 1997. 8, 10
 37. Kresimir Matkovic and Lazlo Neumann. Interactive calibration of the mapping of global illumination values to display devices. In *Proceedings of the Twelfth Spring Conference on Computer Graphics*, Comenius University, Bratislava, Slovakia, June 1996. Available from <http://cg.tuwien.ac.at/wp/SCCG96-proceedings>.
 38. Ann McNamara. Visual perception in realistic image synthesis. *Computer Graphics Forum*, 20(4):211–224, 2001. ISSN 1067-7055. 2, 8, 9, 12, 14
 39. Gary W. Meyer. Wavelength selection for synthetic image generation. *Computer Vision, Graphics, and Image Processing*, 41(1):57–79, January 1988. 19
 40. M. Minnaert. *Light and Color in the Open Air*. Dover, 1954. 6
 41. Karl Mütze, Leonhard Foitzik, Wolfgang Krug, and Günter Schreiber. *ABC der Optik*. VEB Edition, Leipzig, DDR, 1961. 6
 42. Eihachiro Nakamae, Kazufumi Kaneda, Takashi Okamoto, and Tomoyuki Nishita. A lighting model aiming at drive simulators. In Forest Baskett, editor, *Computer Graphics (SIGGRAPH '90 Proceedings)*, volume 24, pages 395–404, August 1990. 12
 43. Kurt Nassau, editor. *Color for Science, Art and Technology*. Elsevier, 1998. 3, 4
 44. L. Neumann, K. Matkovic, and W. Purgathofer. Automatic exposure in computer graphics based on the minimum information loss principle. In *Computer Graphics International 1998*, Hannover, Germany, June 1998. IEEE Computer Society.
 45. László Neumann, Kresimir Matkovic, Attila Neumann, and Werner Purgathofer. Incident light metering in computer graphics. *Computer Graphics Forum*, 17(4):235–247, 1998. ISSN 1067-7055.
 46. Edward D. Palik. *Handbook of Optical Constants of Solids*. Academic Press, 1985. 5
 47. G. B. Parrent and P. Roman. On the matrix formulation of the theory of partial polarization in terms of observables. *Il Nuovo Cimento (English version)*, 15(3):370–388, February 1960. 16
 48. 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. In *Proceedings of SIGGRAPH 98*, Computer Graphics Proceedings, Annual Conference Series, pages 287–298, Orlando, Florida, July 1998. ACM SIGGRAPH / Addison Wesley. ISBN 0-89791-999-8. 2, 10, 12, 15
 49. 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. In Michael Cohen, editor, *SIGGRAPH 98 Conference Proceedings*, Annual Conference Series, pages 287–298. ACM SIGGRAPH, Addison Wesley, July 1998. ISBN 0-89791-999-8. 5
 50. Sumanta N. Pattanaik, Jack E. Tumblin, Hector Yee, and Donald P. Greenberg. Time-dependent visual adaptation for realistic image display. In *Proceedings of ACM SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, pages 47–54. ACM

- Press / ACM SIGGRAPH / Addison Wesley Longman, July 2000. ISBN 1-58113-208-5. 2, 10, 14, 15
51. Mark S. Peercy. Linear color representations for full spectral rendering. In James T. Kajiya, editor, *Computer Graphics (SIGGRAPH '93 Proceedings)*, volume 27, pages 191–198, August 1993. 4, 18
 52. Mark S. Peercy, Benjamin M. Zhu, and Daniel R. Baum. Interactive full spectral rendering. In Pat Hanrahan and Jim Winget, editors, *1995 Symposium on Interactive 3D Graphics*, pages 67–68. ACM SIGGRAPH, April 1995. ISBN 0-89791-736-7. 18
 53. Polarization resource website. Contains pages on viewing polarization with the naked eye. <http://www.polarization.net/>. 6
 54. C. Poynton. Frequently asked questions about gamma. <http://www.inforamp.net/poynton/>. 2, 3
 55. C. Poynton. *A Technical Introduction to Digital Video*. John Wiley and Sons, 1996. 2
 56. A. J. Preetham, Peter Shirley, and Brian E. Smits. A practical analytic model for daylight. In Alyn Rockwood, editor, *Siggraph 1999, Computer Graphics Proceedings, Annual Conference Series*, pages 91–100, Los Angeles, 1999. ACM Siggraph, Addison Wesley Longman. 6, 18
 57. Simon Premoze and Michael Ashikhmin. Rendering natural waters. *Computer Graphics Forum*, 20(4), December 2001. 18
 58. Peter Pringsheim. *Fluoreszenz und Phosphoreszenz im Lichte der neueren Atomtheorie*. Springer, Berlin, 1921. 7
 59. Maria Raso and Alain Fournier. A piecewise polynomial approach to shading using spectral distributions. In *Proceedings of Graphics Interface '91*, pages 40–46, June 1991. 4, 18
 60. Erik Reinhard, Michael Stark, Peter Shirley, and Jim Ferwerda. Photographic tone reproduction for digital images. In *Proceedings of ACM SIGGRAPH 2002, Computer Graphics Proceedings, Annual Conference Series*. ACM Press / ACM SIGGRAPH, July 2002. 10, 13, 15
 61. G. Rougeron and B. Péroche. Color fidelity in computer graphics: A survey. *Computer Graphics Forum*, 17(1):3–16, 1998. ISSN 1067-7055.
 62. Gilles Rougeron and Bernard Péroche. An adaptive representation of spectral data for reflectance computations. In Julie Dorsey and Philipp Slusallek, editors, *Eurographics Rendering Workshop 1997*, pages 127–138, New York City, NY, June 1997. Eurographics, Springer Wien. ISBN 3-211-83001-4. 4, 19
 63. A. Scheel, M. Stamminger, and Hans-Peter Seidel. Tone reproduction for interactive walkthroughs. *Computer Graphics Forum*, 19(3):301–312, August 2000. ISSN 1067-7055. 10, 11, 14, 15
 64. C. Schlick. Quantization techniques for visualization of high dynamic range pictures. In *5th Eurographics Workshop on Rendering*. Eurographics, June 1994. 10, 11, 12, 15
 65. Christophe Schlick. High dynamic range pixels. In *Graphics Gems IV*, pages 422–429. Academic Press, Boston, 1994. ISBN 0-12-336155-9.
 66. Schott Glaswerke corporate website. The interactive technical glass catalog with dispersion data is to be found in the downloads section. <http://www.schott.com/>. 5
 67. Peter Shirley. *Realistic Ray Tracing*. A K Peters, Natick, Massachusetts, 2000. 3, 4
 68. John B. Shumaker. Distribution of optical radiation with respect to polarization. In Fred E. Nicodemus, editor, *Self-Study Manual on Optical Radiation Measurements, Part 1: Concepts*. Optical Physics Division, Institute for Basic Standards, National Bureau of Standards, Washington, D.C., June 1977. 6, 16
 69. Robert Siegel and John R. Howell. *Thermal Radiation Heat Transfer, 3rd Edition*. Hemisphere Publishing Corporation, New York, NY, 1992. 7
 70. B. E. Smits and G. M. Meyer. Newton colors: Simulating interference phenomena in realistic image synthesis. In *Eurographics Workshop on Photosimulation, Realism and Physics in Computer Graphics*, 1989.
 71. Brian Smits. An rgb to spectrum conversion for reflectances. *Journal of Graphics Tools*, 4(4):11–22, 1999. 4
 72. Greg Spencer, Peter S. Shirley, Kurt Zimmerman, and Donald P. Greenberg. Physically-based glare effects for digital images. In *Proceedings of SIGGRAPH 95, Computer Graphics Proceedings, Annual Conference Series*, pages 325–334, Los Angeles, California, August 1995. ACM SIGGRAPH / Addison Wesley. ISBN 0-201-84776-0. 10, 12, 15
 73. Jos Stam. Diffraction shaders. In Alyn Rockwood, editor, *Siggraph 1999, Computer Graphics Proceedings, Annual Conference Series*, pages 101–110, Los Angeles, 1999. ACM Siggraph, Addison Wesley Longman.
 74. S. S. Stevens and J. C. Stevens. Brightness function: Parametric effects of adaptation and contrast. *Journal of the Optical Society of America*, 53(1139), 1960. 9
 75. Y. Sun, M. Drew, and F. Fracchia. Representing spectral functions by a composite model of smooth and spiky components for efficient full-spectrum photorealism. In

- Proceedings of the Workshop on Photometric Modeling for Computer Vision and Graphics (PMCVG-99)*, pages 4–11, Los Alamitos, June 22 1999. IEEE. 18
76. Yinlong Sun, F. David Fracchia, Thomas W. Calvert, and Mark S. Drew. Deriving spectra from colors and rendering light interference. *IEEE Computer Graphics and Applications*, 19(4):61–67, July/August 1999. 18
 77. Yinlong Sun, F. David Fracchia, and Mark S. Drew. A composite model for representing spectral functions. Technical Report TR 1998-18, School of Computing Science, Simon Fraser University, Burnaby, BC, Canada, November 1998.
 78. Yinlong Sun, F. David Fracchia, Mark S. Drew, and Thomas W. Calvert. A spectrally based framework for realistic image synthesis. In *The Visual Computer*, volume 17(7), pages 429–444. Springer, 2001. 4, 18
 79. David C. Tannenbaum, Peter Tannenbaum, and Michael J. Wozny. Polarization and birefringency considerations in rendering. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pages 221–222. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0. 14, 16
 80. Spencer W. Thomas. Dispersive refraction in ray tracing. *The Visual Computer*, 2(1):3–8, January 1986. 5
 81. J. Tumblin. *Three Methods of Detail-Preserving Contrast Reduction for Displayed Images*. Phd thesis, Georgia Institute of Technology, December 1999. 2, 3, 12, 13
 82. Jack Tumblin, Jessica K. Hodgins, and Brian K. Guenter. Two methods for display of high contrast images. *ACM Transactions on Graphics*, 18(1):56–94, January 1999. ISSN 0730-0301. 10, 11, 12, 13, 14, 15
 83. Jack Tumblin and Holly E. Rushmeier. Tone reproduction for realistic images. *IEEE Computer Graphics & Applications*, 13(6):42–48, November 1993. 9, 10, 15
 84. Jack Tumblin and Greg Turk. Lcis: A boundary hierarchy for detail-preserving contrast reduction. In *Proceedings of SIGGRAPH 99*, Computer Graphics Proceedings, Annual Conference Series, pages 83–90, Los Angeles, California, August 1999. ACM SIGGRAPH / Addison Wesley Longman. ISBN 0-20148-560-5. 10, 13, 15
 85. Greg Ward. A contrast-based scalefactor for luminance display. In *Graphics Gems IV*, pages 415–421. Academic Press, Boston, 1994. ISBN 0-12-336155-9. 9, 10, 11, 12, 15
 86. Gregory Ward. High dynamic range imaging. In *Proceedings of the Ninth Colour Imaging Conference*, November 2001. 19
 87. Gregory J. Ward. The RADIANCE lighting simulation and rendering system. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pages 459–472. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0. 18
 88. Turner Whitted. An improved illumination model for shaded display. *IEEE Computer Graphics & Applications*, 23(6):343–349, June 1980.
 89. Alexander Wilkie, Robert F. Tobler, and Werner Purgathofer. Raytracing of dispersion effects in transparent materials. In *WSCG 2000 Conference Proceedings*, 2000. 5
 90. Alexander Wilkie, Robert F. Tobler, and Werner Purgathofer. Combined rendering of polarization and fluorescence effects. In S. J. Gortler and K. Myszkowski, editors, *Proceedings of the 12th Eurographics Workshop on Rendering*, pages 197–204, London, UK, June 25–27 2001. 14, 16, 17
 91. Lawrence B. Wolff and David Kurlander. Ray tracing with polarization parameters. *IEEE Computer Graphics & Applications*, 10(6):44–55, November 1990. 7, 14
 92. G. Wyszecki and W. S. Stiles. Color science: Concepts and methods, quantitative data and formulae. *John Wiley and Sons, New York*, ISBN 0-471-02106-7, 1982. 3, 4
 93. Hector Yee, Sumanta Pattanaik, and Donald P. Greenberg. Spatiotemporal sensitivity and visual attention for efficient rendering of dynamic environments. *ACM Transactions on Graphics*, 20(1):39–65, January 2001. ISSN 0730-0301.
 94. Ying Yuan, Tosiyasu L. Kunii, Naota Inamoto, and Lining Sun. Gemstonefire: adaptive dispersive ray tracing of polygons. *The Visual Computer*, 4(5):259–270, November 1988. 5