

A Practical Approach to Spectral Volume Rendering

Steven Bergner, Torsten Möller, Melanie Tory, Mark S. Drew
 Graphics, Usability, and Visualization Lab
 Simon Fraser University

Abstract—To make a spectral representation of colour practicable for volume rendering, a new low-dimension subspace method is used to act as the carrier of spectral information. With that model spectral light material interaction can be integrated into existing volume rendering methods at almost no penalty. In addition, slow rendering methods can profit from the new technique of post-illumination — generating spectral images in real-time for arbitrary light spectra under a fixed viewpoint. Thus, the capability of spectral rendering to create distinct impressions of a scene under different lighting conditions is established as a method of real-time interaction. Although we use an achromatic opacity in our rendering, we show how spectral rendering permits different dataset features to be emphasized or hidden as long as they have not been entirely obscured. The use of post-illumination is an order of magnitude faster than changing the transfer function and repeating the projection step. To put the user in control of the spectral visualization we devise a new widget, a “light-dial”, for interactively changing the illumination and include a usability study of this new light space exploration tool. Applied to spectral transfer functions, different lights bring out or hide specific qualities of the data. In conjunction with post-illumination this provides a new means for preparing data for visualization and forms a new degree of freedom for guided exploration of volumetric data sets.

I. INTRODUCTION

DIRECT volume rendering allows views beneath surfaces, showing features inside volumetric data sets. To make certain parts of the scene visible, a transfer function assigns distinct colours and opacities (alpha) to the data points. Our approach is to use spectral reflectances instead

of RGB values to set up the transfer function. Spectral materials in the scene now appear as different colours depending on the illuminating light (although alpha compositing is not affected). This changes the appearance of the scene every time a new light is set. Under this framework, materials and lights can be specially designed to achieve certain visual effects. Since we use achromatic absorption, an object which is obscured by large regions of high opacity will not be made visible by re-lighting. It has been our experience, however, that the extra dimensions of information enabled by spectral rendering decrease the importance of opacity assignment and occlusion as the fundamental means of determining feature visibility. Changing the illumination of a scene thus becomes a new means for the user to interact with the visualization. In order to deploy this strategy as a useful tool for exploration a new user interface is proposed.

Many algorithms gain speed by assuming that the transfer function is fixed. While this seems a rather restrictive assumption, we demonstrate in this paper that we can still manipulate the visualization to potentially gain more insight into our data set simply by controlling the light sources. We show how to do this in real-time for any volume rendering algorithm. Extending the transfer function mapping to non-restrictive spectra enables us to exploit to our advantage certain phenomena of colour that are usually undesired. Our main idea is to make use of the concept of *metamers*. Differing light spectra that appear perceptually identical under a certain light are termed metameric. This means that such spectra differ only by a vector in the null space of the functions transforming spectra into colours (e.g., the spectral curves of the R, G, and B sensors of a virtual camera). By assigning metameric colours to certain data intensity ranges (using a *spectral transfer function*), we have the ability to visually

merge these ranges into one colour under a specific light. By changing the light spectra we can visually separate these ranges. This has the potential to help the user explore the relationships between certain intensity intervals or materials (see §IV-A for more details).

In the following we distinguish between light, reflectances and colours. Colour is a psychophysical quantity of light that is sensed by the human eye. As depicted in Fig. 1, light is sent out by sources in the environment and eventually reflected or refracted by other objects until it arrives at the eye. For graphics, we can use the *spectral power distribution* (SPD) of the electro-magnetic (EM) wave – the wavelength-dependent magnitude of the EM field’s Poynting vector [20] – with visible wavelength range 400 nm to 700 nm, instead of full optical physics. When talking about physical correctness




Fig. 1. Light may reflect off many surfaces before it reaches the eye where it is perceived as “colour”.

and error we have to keep in mind that for visualization, expressiveness of images comes before photo-realism or physical accuracy. Furthermore, there are typically data sets, such as CT or MRI, that do not contain sufficient information on the original object. However, spectra can be used to great benefit in visualization. In particular, their additional dimensionality allows for more versatile image manipulation. This is the central idea behind the interaction we propose.

Spectral colour models aspire to preserve the necessary physical properties of spectra to allow illumination calculations without *perceivable* errors. To minimize computational costs we use a new low-dimensional linear representation called the *spectral factor model* of colour described in §III.

Our new technique of post-illumination described in §V allows one to efficiently re-render a scene

for changing light sources. For the interactive manipulation of a spectral rendition via our post-illumination scheme a new “light-dial” widget is devised in §IV-B. This paper is an extension of our previous work [2]. Here, we (1) show details on the spectral factor model (§III), (2) extend various rendering algorithms to incorporate post-illumination (§VI), and (3) present a usability study including a keystroke-level analysis as well as a usability inspection of the new interface (see §VII).

II. RELATED WORK

A. Volume Rendering

The past five years have seen great advances in the area of real-time volume graphics. Much-improved consumer hardware (referred to as GPU) is available that has been efficiently exploited for volume rendering using texture mapping capabilities as well as many new features of modern graphics cards [6]. Other improved data structures such as an adjacency list and an RLE based data structure were introduced to speed up the splatting algorithm tremendously [11], [17]. Fourier Volume Rendering [14], [21] enables software-only real-time rendering (with quality degradation). Our method of image-based spectral re-lighting can be incorporated into existing volume rendering methods as discussed later in §VI. A performance comparison shows that the speed of projection for arbitrary viewpoint is only slightly compromised gaining interactive light manipulation.

To change the colour scale in the rendered image Kaneda et al. [10] linearly combine renditions of different Fourier basis transfer functions. This requires multiple rendering passes. Their approach is restricted to re-colouring pseudo-colour images, while we seek to combine correct spectral light-material interaction with fast re-rendering.

Finally, Noordmans et al. [16] have integrated a spectral colour model into the volume rendering pipeline. To speed up their rendering, they go from rendering full/actual spectra to rendering coefficients of materials that are assumed to occupy distinct spectral bands. The difference in our approach is to gain performance by using an optimized linear colour model throughout the entire rendering process. This allows us to perform full spectral multiplications with minimal computational effort. The rendering is not restricted to a given number of

materials, as for Noordmans et al. Our method also provides more freedom in the way spectra change while interacting within the scene. This is essential for more advanced illumination models. Further, we discuss guidelines, not addressed by Noordmans et al., of how to sensibly create and use artificial spectra, an important issue for volume rendering and for computer graphics in general.

B. Representing Light

The most straightforward method of dealing with spectral information is to simply carry spectrally point-sampled data through any rendering calculation. Sufficiently sampling at a uniform interval, however, would mean too many samples¹.

Linear models express an SPD as a weighted sum of basis functions. They have typically been based on eigenfunction expansions of collections of spectral power distributions. Such finite-dimensional models offer optimal representation, in terms of variance-accounted-for, and generally have low dimension (6 to 8). Peercy [18] made good use of such models in an attempt to integrate spectra into surface graphics. However, one is still left with a matrix multiply at each interaction with matter and this presents a stumbling block for using the method.

III. SPECTRAL FACTOR MODEL

We present a new colour-vector representation that adopts the strength of linear models while avoiding spectral multiplication: light interaction is reduced to a simple extension of RGB-colour componentwise multiplication — a *spectral factor model* [5].

The linear colour model we use here is based on an extension of the idea of using *spectral sharpening* in colour constancy algorithms in computer vision [8] to spectral bases. In graphics we usually multiply RGB componentwise to form colour. In surface graphics, for example, illuminant RGB is multiplied times surface RGB in order to generate a product RGB. When lighting changes, we would like to model RGB change as a simple diagonal transform. This is not physically accurate,

¹Real surfaces and illuminants are typically measured at more than 100 sample points and represented using 31 samples from 400 nm to 700 nm at 10 nm intervals.

but can be made more so by a preliminary matrix transformation of camera sensors to generate an intermediary colour space. This transform is in fact included in many digital cameras, and is called “spectral sharpening”, since camera sensors are thereby made effectively more narrowband. For multiple light-surface interactions, however, such a simple transform will not lead to enough accuracy and we must fall back on full spectra.

Here we can apply the same idea to the linear model basis set: we form combinations of the basis set vectors that are optimally narrowband. That is, we pre-‘sharpen’ the basis by a simple matrix transform and then agree to operate within the sharpened basis for all surface or volume interactions [5]. Note that no information is lost by such a transform, and accuracy to within the adopted dimensionality of the underlying finite-dimensional model is maintained.

Then indeed we can represent spectral interactions in terms of the low-dimensional coefficients (typically 5 to 7) and calculate interactions using componentwise multiplication of the coefficients. Using the new basis, the multiplication of spectral functions in the linear subspace is reduced to a simple componentwise multiply involving the coefficients ε_i of the current light, $E(\lambda)$, times coefficients σ_i for the next interaction surface, reflectance or transmission function $S(\lambda)$. If

$$C(\lambda) = E(\lambda)S(\lambda) \quad (1)$$

then, projected into the d -dimensional linear subspace, we have approximately

$$\gamma_i \simeq \varepsilon_i \sigma_i, \quad i = 1..d \quad . \quad (2)$$

where γ_i are the basis coefficients for the colour signal $C(\lambda)$. Note that instead we could simply be using sums of delta function spectra similarly, but then would be giving up the power of the basis set to accurately model actual spectra. Here we retain simplicity and low dimensionality without abandoning physical spectra.

IV. INTERACTION ICON

While material assignment has been used successfully in volume rendering before [4], using full spectra provides a good deal more freedom in designing transfer functions. A careful strategy is required to sensibly make use of this freedom. A




Fig. 2. **Spectral raycasting** of the engine data set under three different lights. The parts of the engine have three different materials assigned. The images were generated using post-illumination — no re-rendering took place, just an image-based post-processing.

user-oriented integration of spectral concepts into a volume rendering pipeline will have to consider the following steps:

- 1) Materials and light sources have to be chosen or *designed* with close attention to their use in volume rendering: both metamers, with $C_1(\lambda)$ and $C_2(\lambda)$ yielding the same RGB, can and should be used to pick out particular regions of interest.
- 2) Materials are assigned to spectral transfer function values in order to create distinct perceptual discrimination.
- 3) Finally, we are able to navigate through different visualizations by controlling the light source. Since multiple dimensions (n weights for n light spectra) are involved we create a new interaction icon.

While the second step is a straightforward material based transfer function design, steps one and three deserve further discussion.

A. Material Design

By assigning distinct reflectances we are able to separate the appearance of materials. However, we also have *metameric materials* that under a particular light source result in colours that are identical for our visual sensors. Hence, we have the ability to merge certain materials and therefore guide the user’s attention towards or away from these materials. Fig. 2 illustrates this effect. Under light 1 (Fig. 2a) both materials of the block look the same. The entire engine is perceived

as a homogeneous structure. Blending to light 2 (Fig. 2b), the metamer effect breaks down and the materials become distinguishable. This process of newly emerging details is likely to be noticed by the user.

Another principle in spectral material design is to make use of the effect of *colour constancy*. In Fig. 2 we have created light sources, in connection with materials, that will leave the colour of one material the same and change only the colour of the materials we intend to influence with that light — the engine block keeps its green colour under both lights.

Since traditional colour theory usually describes only phenomena that occur in the real world, this restricts us to non-artificial metameric materials. However, it would be desirable to design materials that disappear entirely from the image, if so desired by the user. In colour science, these materials are called *metameric blacks* [23], i.e., materials with zero RGB under some specific set of lights. Such a material is assigned to the surrounding reconstruction noise in Fig. 2. This ‘smoke’ becomes visible as we blend over to light 3 in Fig. 2c. In the other images it is rendered as well, but under the first two lights it remains black. In rendering, alpha compositing still takes place, so occlusion of other materials is in fact not undone. Nevertheless, it is possible to introduce materials that only scatter light but do not actually have an alpha value to diminish light that comes from the back. If those X-ray materials turn black under a certain light they entirely disappear from the image sum and free the view to things that they obscure. Since they will

always have an X-ray like look, it is not possible to draw crisp surfaces with such materials. This idea is applied in our implementation of spectral Fourier volume rendering discussed in section VI.

While metamers can be used to reduce complexity in images, constant colours are useful when blending from one light source to another. Having some things change their appearance while others remain the same is likely to grab the observer’s attention. Thus, such materials can be used to guide the user’s attention to certain aspects of the data during interactive exploration. Within the same visualization (without any change of the transfer function) we are able to emphasize and de-emphasize particular regions of interest in connection with particular light sources. Therefore we can help the user understand the relationships between these materials.

Having an idea of what colour constellations are desirable for visualization, we now have the task of finding spectra for lights and materials that fit these criteria. One approach is to use materials previously acquired from real world scenes. Nevertheless, for most applications it would be impractical to have to look for real materials and lights that fit the desired properties (metamerism, colour constancy, metameric blacks). For that purpose we employ a design procedure [1] that combines real spectra with new artificial but physically-based ones generating a set of lights and reflectances.

B. Navigating Visibility

We can design materials and lights such that each light influences exactly one material. Hence a linear combination of the lights will create a mixed rendition of the scene. Assuming that we have m materials and n lights, we have an n -dimensional space to navigate to produce different visualizations. This is a difficult task for the user.

We developed an interaction metaphor that includes all light sources within one interaction widget, which we call the “light dial”. This is a two-dimensional n -gon slider, mapping a 2D space into the n -dimensional parameter space (see Fig. 3). Our n light sources are the vertices of the n -gon. Any position on the plane characterizes a weighted sum of the light sources that make up the n -gon. This way, the higher-dimensional parameter space is mapped to a 2D topology of nodes.




Fig. 3. **The light dial** – interface to control the mixture of lights using normalized inverse distances of the mixture selector (yellow circle) to the light nodes (bulb icons) (see Eq. 3).

Each control node has a position on the screen and represents a light source, and the mouse is used to freely move the mixture selector over and between the nodes (e.g., the yellow dot position in Fig. 3). Position is used to determine a scalar weight for each node, with a value that grows as the selector comes closer to a node. We can compute a weight $L_i(\mathbf{x})$ applied to the light for a selector position \mathbf{x} in the following manner:

$$L_i(\mathbf{x}) = \prod_{k=1, k \neq i}^n \frac{\|\mathbf{x} - \mathbf{x}_k\|}{\|\mathbf{x}_i - \mathbf{x}_k\|}. \quad (3)$$

To make the weights usable for combining light sources, we normalize the sum of all weights L_i to one and scale it with a separate intensity slider. When the position of the dot coincides with the position of a light bulb in the widget, the influence of all other light sources is eliminated. The influence of a source can be entirely removed by switching it off. Hence, the light dial allows one to focus on the influence of a single light to the materials used in the scene. It has not been designed to navigate the full n -D space at once. Instead, the intention is to conveniently slide from one pure light to the next pure light.² Additionally, it is possible to correlate dimensions (making lights have similar intensity) by moving two light bulbs closer to each other.

²Nevertheless, it is still possible to navigate the entire space of weight combinations. Imagine a concentric arrangement of the light bulbs around the light dot in the middle: all lights have the same weights. Now, a light bulb may be dragged instead of the dot, moved closer to or farther away from the center (where the dot is). Since only one distance is changing, it is possible to change only one weight. To actually have just one weight changing, the separate overall intensity slider would have to be moved correspondingly.

This modifies the shape of the 2D mapping of the parameter space and allows the user to reach new parameter constellations.

To summarize, our light dial implements three different interaction metaphors:

- 1) move yellow dot (change weights by distance) — the object is weighted towards the lights we want to see it under,
- 2) move light nodes — group lights, or drag lights away to lower their influence (mathematically speaking: change shape of 2D mapping of higher-dimensional parameter space, correlate dimensions by decreasing spatial distance in widget plane), and
- 3) switch lights on/off (reducing dimensionality, comparable to taking a light node and moving it far away; deactivation of the light allows it to keep its position in the mapping plane.)

This makes the 2D dial a convenient interface to navigate through the colour schemes. The light dial metaphor may be useful for other tasks, such as navigating between different points of view, controlling alpha values of different segments, etc. A more thorough evaluation of this widget is provided in §VII.

V. POST-ILLUMINATION

By fading between different lights the whole scene changes its appearance. To make this technique usable for data exploration, we propose the method of *post-illumination*. It allows to re-render the scene at interactive frame-rates by separating shading and compositing from the actual computation of an RGB colour. Fig. 4 illustrates the notation that is adopted in this section. Looking at the rendering integral as discussed by Max [15] it is possible to factor out the light source yielding the following equation:

$$I_i(\mathbf{x}, \boldsymbol{\omega}_i, \lambda) = L(\lambda) \cdot \tilde{I}_i(\mathbf{x}, \boldsymbol{\omega}_i, \lambda), \quad (4)$$

It describes the irradiance I_i depending on the wavelength λ and the position \mathbf{x} on a ray of length D from direction $\boldsymbol{\omega}_i$. In this model the point \mathbf{x} can either be a voxel in the volume or a pixel on the projection screen at which we collect the incoming light. The second term, \tilde{I}_i , is a ratio indicating which portion of light L is transferred through the




Fig. 4. Illustration of symbols and abbreviations used in §V.

scene. It is calculated as

$$\tilde{I}_i(\mathbf{x}, \boldsymbol{\omega}_i, \lambda) = \tilde{I}_i(\mathbf{d}, \boldsymbol{\omega}_i, \lambda)T(\mathbf{x}, \boldsymbol{\omega}_i, \lambda, D) + \int_0^D \tilde{I}_r(\mathbf{x} - s\boldsymbol{\omega}_i, \boldsymbol{\omega}_i, \lambda) \cdot T(\mathbf{x}, \boldsymbol{\omega}_i, \lambda, s) ds, \quad (5)$$

where $T(\mathbf{x}, \boldsymbol{\omega}, \lambda, s) = e^{-\int_0^s \tau(\mathbf{x} - t\boldsymbol{\omega}, \lambda) dt}$ is the extinction integral that accounts for the effect of absorption τ . The first term is the irradiance from outside the volume at a point $\mathbf{d} = \mathbf{x} - D\boldsymbol{\omega}_i$. Since the light source $L(\lambda)$ is factored out from the integral, the radiance I_r is modified to become a ratio \tilde{I}_r . This ratio can be understood as the amount of irradiance that is scattered from direction $\boldsymbol{\omega}_r$. More precisely, it is formed as

$$\tilde{I}_r(\mathbf{x}, \boldsymbol{\omega}_r, \lambda) = \int_{4\pi} \rho(\mathbf{x}, \boldsymbol{\omega}_r, \boldsymbol{\omega}_i, \lambda) \cdot \left(\tilde{I}_i(\mathbf{x}, \boldsymbol{\omega}_i, \lambda) + \tilde{I}_i^L(\mathbf{x}, \boldsymbol{\omega}_i, \lambda) \right) d\boldsymbol{\omega}_i, \quad (6)$$

containing the product of the bi-directional reflection distribution function ρ (BRDF) and the irradiance integrated over all unit directions $\boldsymbol{\omega}_i$. Here, the shadowed amount of direct light from source L is captured by \tilde{I}_i^L and the indirectly scattered light is the recursive term \tilde{I}_i from Eq. 5.

Our implementation uses raycasting with a local illumination model. This simplifies \tilde{I}_r^L to just become the reflectance of the local material, weighted by the Phong shading coefficients (ambient, diffuse and specular). Even anisotropic and iridescent effects could be incorporated here as a wavelength dependent weighting. As for the local illumination,

\tilde{I}_i is zero for all directions other than along the ray $\omega_i = \omega_r$. Furthermore, our shadowing ratio \tilde{I}_i^L is constant 1 for all voxels and directions (no shadowing). An example rendition incorporating the use of specular lighting is shown in Fig. 2, where we can observe specular highlights at the solid inner rings inside the engine. The images are generated using an achromatic absorption (alpha blending). Nevertheless, the use of metamers and metameric blacks enables us to influence the visibility of additional structures even after the spectral image has been generated by simply changing the illuminating light spectrum.

The postponed multiplication of the light source allows one to generate new images for different light spectra without repeating the projection and compositing. After all light bouncing and transmission is complete, only one matrix multiplication in the linear colour model per final image pixel is left to calculate the updated RGB image for the scene. This gives the necessary speedup for interactive re-illumination.

VI. SPECTRA IN VOLUME RENDERING TECHNIQUES

Many benefits accrue to different rendering methods from incorporating a spectral colour model. Here we explain how several classic techniques (a non-exhaustive list) are impacted by spectral extensions.

A. Raycasting

This technique is a straightforward implementation of a simplified version of the rendering integral discussed in §V. Raycasting provides high rendering quality, but due to random memory access it is not very fast in its standard implementation. This technique derives the biggest profit from using post-illumination. Rendering reflectances to the screen takes only slightly longer than RGB rendering [2]. Once the spectral image is generated, it can be illuminated with different light sources, via a 3×7 matrix transformation per pixel, combining the illumination matrix and the colour transform to produce RGB. The examples previously shown in Fig. 2 illustrate the visual effect of post-illumination applied to the raycasted dataset of an engine.




Fig. 5. **Spectral Splatting** of hipiph data set. The lower intensities are assigned a low-alpha material (a) glowing white under light 1, or (b) being black under light 2.

B. Splatting

Splatting is typically a very fast volume rendering method, introduced by Westover [22]. Current graphics cards are based on RGB compositing. This is the major constraint when using spectral rendering. In order to combine fast splatting with spectral light material interaction we have to limit the rendering to flat absorption (alpha blending). There are two ways to build a spectral splatter. One is to perform a multi-pass rendering until all channels of the colour model have been projected to the screen (e.g. two passes for a 6 channel factor model). The resulting image could then be post-illuminated just as discussed above. An alternative approach is to simply use an RGB transfer function whose colours are made from spectral reflectances and the active light source. If the colour (RGB) for a voxel is determined, it can be splatted to the screen using standard graphics hardware.

One way to further improve performance is to use vertex programs to do the illumination calculation at each voxel with spectral light sources and perform the colour transformation to RGB. The colour of the light source, its position, and the colour transformation matrix (size $3 \times (7 + 1)$) are stored in the vertex program parameter space. The 7-component voxel reflectance is passed utilizing the primary and secondary colour registers. The conversion from the 7-component model can be done by splitting the conversion matrix horizontally into two 4×4 matrices (last column is zero). These are applied to the halves of the reflected light colour vector accordingly. By adding the results the final colour is obtained. Such operations are basic in vertex programs and can be done in a few lines of code. Using the spectral splatter, the $64 \times 64 \times 64$




Fig. 6. **Spectral Fourier Volume Rendering** of the UNC-Brain data set (CT-scan). Each light illuminates one material keeping the other two black. The images show three different mixtures of lights illuminating: (a) interior only, (b) three different regions, (c) frontal skull. As to the x-ray character of the images the black material entirely disappears from the image (no occlusion).

hipiph dataset in Fig. 5 can be rendered at 18 frames per second still allowing for an interactive manipulation of the transfer function and spectral light colour. A similar RGB splatter without spectra is only slightly faster as it can be seen in Table I.

C. Fourier Volume Rendering

The basic idea of Fourier volume rendering (FVR), as introduced by Malzbender [14], is to calculate the projection of a data set to the screen by extracting a slice of the volume in the frequency domain. This makes FVR a very powerful technique, especially for large data sets. Because the intensities are simply added up the result is an x-ray like image. Totsuka and Levoy [21] further improved the technique by incorporating hemispherical shading and depth-cueing. Recent extensions by Entezari et al. even allow for specular highlights [7].

To integrate spectral illumination into FVR our approach is to render the 7 components of the linear colour model separately. Post-illumination can be carried out by projecting only the reflectances and performing the illumination calculation (multiplying the light spectrum) afterwards in the image. Examples are shown in Fig. 6 and Fig. 7. The first figure makes extensive use of metameric blacks to toggle visibility of different segments. Spectral materials are assigned such that metamers make different segments merge together; and segments coloured with metameric blacks entirely disappear from the image sum for a certain light source. FVR is a fast technique for pure software implementation. Using spectra and post-illumination in FVR, allows re-interpreting the spectra in the transfer

function in a fast image-based operation and makes an expensive rerun of the pre-processing unnecessary. Specific parts of the data can be separated, merged, or faded out without changing the transfer function. This not only improves the visual quality of the renderings but also adds a new degree of interaction to FVR.




Fig. 7. **Spectral Fourier Volume Rendering** of a tomato (MRI) rendered with (a) two metameric materials, and (b) the same materials being non-metameric under a different light.

D. Summary

An evaluation of the performance of different spectral rendering techniques is provided in Table I. The techniques are from left to right: spectral raycasting and separate post-illumination (PI) in comparison with plain RGB raycasting, spectral splatting vs. RGB splatting, and spectral Fourier volume rendering (SFVR). All timings were taken on a 2.4 GHz Pentium 4 with a GeForce 4 using a constant image size of 512×512 . The time is given in frames per second for projecting a frame / re-illuminating a frame. The last column also

Dataset	Fill	Spec-RC	PI-RC	RGB-RC	Spec-Splat	RGB-Splat	Spec-FVR/PI-FVR
hipiph	97%	0.22	9.2	0.24	18.5	22.2	50/166 (1s)
tomato	72%	0.15	11.6	0.17	0.78	1.7	5.8/34.5 (85s)
uncbrain	62%	0.15	13.3	0.17	0.3	0.4	15.8/30.3 (85s)
engine	79%	0.13	10.7	0.15	0.5	0.6	8.7/47.6 (85s)
frog	69%	0.08	12.0	0.09	0.31	0.38	(<i>n/a</i>)

TABLE I

PERFORMANCE OF DIFFERENT SPECTRAL RENDERING TECHNIQUES (IN FPS), COMPUTING IMAGES FROM A NEW VIEWPOINT. *Fill* DENOTES THE AMOUNT OF NON-BLACK SCREEN PIXELS. *PI* IS POST-ILLUMINATION (FOR FIXED VIEWPOINT) ON THE SPECTRAL IMAGE. DATASET DIMENSIONS: HIPIPH 64^3 , TOMATO $256^2 \times 64$, UNCBRAIN $256^2 \times 145$, ENGINE $256^2 \times 128$, FROG $500 \times 470 \times 138$

contains time for preprocessing for SFVR. PI is independent of the rendering technique. Nevertheless, the timings are given for images from the spectral raycasting. The method has been optimized to only consider pixels of non-zero reflectance. Thus, the screen *fill* correlates to the PI timings. Note that PI is an image-based operation and not necessarily faster than object based methods, as it is the case for the spectral splatting of *hipiph*. As discussed above, our variant of the spectral splatter recomputes the projection for re-illumination, because compositing is done in RGB. Note that the images of SFVR internally do not actually have the size of the display but instead are obtained from the extracted slice from the frequency cube. Thus, the timing of PI on SFVR is not directly comparable to the spectral raycasting. The two major conclusions that can be drawn from the table are: 1.) the performance of spectral rendering compared with corresponding RGB techniques (e.g., spectral splatting) shows very competitive framerates, and 2.) the use of post-illumination is an order of magnitude faster than changing the transfer function and repeating the projection step.

VII. USABILITY ASSESSMENT OF THE LIGHT DIAL

To utilize spectral volume rendering, users would typically (1) set up materials and lights prior to a visualization session and then (2) examine data by setting up transfer functions and manipulating lights to change visibility (as established in section IV). Designing a good interface for setting up materials and lights is complex; our interface has the appropriate functionality but non-ideal usability. We leave development of a good material/light designer for future work. Designing good transfer function

editors is also challenging, but is not unique to spectral volume rendering. Hence, we focus on evaluating our proposed light-dial for navigating visibility.

We evaluated the light dial and its conventional counterpart — a multiple-slider interface having a separate intensity slider for each light. The following two methods of evaluation were chosen:

- 1) Keystroke-level analysis [3] to compare interface speed, and
- 2) Expert assessment of the interfaces via usability inspection [13].

A. Keystroke-Level Analysis

In keystroke-level analysis, tasks are broken into keystroke-level actions (clicking a key, dragging the mouse, etc.). Time for an expert user to complete a task is predicted by adding together average times for the keystroke-level actions [3].

We made several assumptions in our keystroke-level analysis:

- 1) Some actions (e.g., homing hands on the mouse) are the same for both interfaces and can therefore be ignored.
- 2) Both interfaces are constrained such that the largest mouse movement is 5 cm.
- 3) An average mouse movement is half the largest movement.

We used component actions and times from [3] to predict time required to adjust settings in each interface, and considered several situations (e.g., turning one or all light(s) on). We considered the interfaces with and without shortcuts (i.e., clicking at the desired end location without dragging).

Shortcuts work for any slider in the multiple-sliders interface and the yellow slider in the light dial, but not for light sliders in the light dial since it would be impossible to know which light a user intended to move.

Results favour the light dial when only a few lights are active. With 4 lights or less, the light dial is equal or faster than multiple-sliders for all tasks. With more than four lights, the light dial is as good or better than multiple-sliders when only the yellow slider must be moved (i.e., to turn on one light, all lights, or sets of lights that are close together and far from undesired lights). These results are the same whether shortcuts are allowed or not.

Setting up constellations with the light dial is inefficient because dragging lights takes time. However, as previously mentioned, navigating between constellations with the light dial is efficient. Since materials and lights are designed to match, pure lights (one light source alone) may be most interesting and arbitrary light combinations may be used less often. Furthermore, users who do set up new light dial “shapes” may then spend substantial time moving the main slider within that shape before altering the shape again. In this case, the cost of setting up the shape may be offset by the speed gained during the second navigation phase. However, users who set up arbitrary light combinations quite often are faster with the multiple-slider interface.

B. Usability Study

a) Method: Three usability experts were asked to assess the interfaces together and separately using usability inspection techniques [13]. All three experts were graduate students in human-computer interaction at Simon Fraser University. None had prior knowledge of the spectral volume rendering project. Evaluations were conducted independently.

Lights and materials were set up ahead of time. A 4-light setup was used. Experts were given descriptions of typical end users and typical tasks (data exploration, matching the display to target images, and returning to previous settings). They answered questions about advantages, disadvantages, and problems with the interfaces, and evaluated them with respect to the following heuristics: ease of learning, of use, of exploring a data set, of matching target images, ease of undo, speed of ex-

ploring, speed of reaching target images, and speed of undo. Heuristics were assessed on a 5-point rating scale and other questions were open-ended. After a written report was completed, evaluators discussed their opinions with the experimenter.

b) Usability Inspection Results: Fig. 8 displays average ratings in the heuristic evaluation. Multiple sliders were easier to learn and use (overall), probably because of their familiarity. Exploration and undo were easier with the light dial, probably because a wide range of light levels could be explored by simply moving one slider and relative positions of controls were easier to remember with the 2D representation. In terms of speed, the light dial was rated faster for all tasks, agreeing with the keystroke-level analysis. Having both interaction techniques available side-by-side generally got good ratings and appears to be a good compromise. In this case, a user can choose which interface to use and the other updates automatically.

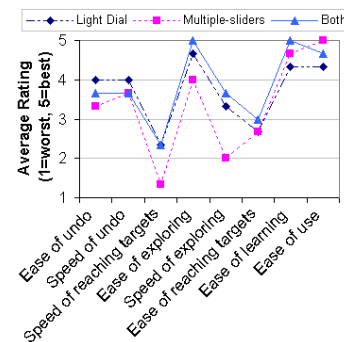


Fig. 8. Average ratings in heuristic evaluation.

All evaluators agreed that both interfaces should be available since each is useful for different tasks. The multiple-sliders interface was considered useful for understanding the contribution of each light and for resetting these contributions exactly during an undo operation. The light dial was useful for quickly exploring, turning lights on and off, and navigating preset constellations. However, judging the levels of each light targets in the light dial was difficult. Some evaluators favoured having both interfaces in the same dialog box, but others suggested that the light dial be the main interface and the multiple-slider interface be accessible as an auxiliary widget. One report predicted that prolonged use of the light dial would cause less wrist strain than the multiple-slider interface. Overall, the usability study indi-

cated that users can perform basic lighting changes via the light dial with little training.

VIII. DISCUSSION AND FUTURE WORK

To construct a spectral transfer function from a set of given materials is just as straightforward as it is to create a traditional RGB-based visualization. The important difference is that the manipulation of the light source allows for more versatile interactions later on. Nevertheless, creating one's own spectra still requires substantial work and a number of deliberate decisions. The actual benefits that justify the additional work can be summarized as follows:

- 1) The scene may bear more information than it is possible to see with three-dimensional RGB,
- 2) the interaction with the scene is more versatile since light changes become a means of exploration,
- 3) the user's attention may be guided by fading between different visualizations that are inherent to the incorporated spectra,
- 4) the visualization may more truthfully reflect the physical behaviour of given materials.

After the projection step is done, different impressions may arise from the visualization by post-illuminating the image with different lights. The user can interact with the image by manipulating the light. Thus, through one spectral scene setup, different messages may be conveyed to the observer. Developing this method up to a state of interactive frame rates for any data and renderer is seen as a major contribution of this paper.

A different method achieving a similar effect is to render the dataset with different traditional (RGB-based) transfer functions and to blend between the resulting images. Despite the fact that this method may require more rendering passes, it is a valid alternative to our approach. The important difference to post-illumination is that we can apply any kind of light spectrum to the unilluminated image. This is not possible when combining prerendered images of RGB-based transfer functions. In our current framework this advantage is not really used. The only situation where this can actually be observed is when designing new lights and seeing their immediate effect on the scene without re-rendering.

This is possible within the framework, but has not been further evaluated as of yet.

The system presented here opens opportunities for different types of applications. One direction would be photo-realistic volume rendering. As our method computes correct colours for light-material interactions, it is possible to incorporate real-world materials into the scene. This could be useful to create more convincing training scenarios for medical applications. For this it might be interesting to incorporate the novel ideas of spectral sub-surface scattering [12].

To date, not much work has been done on interactively rendering spectral BRDFs. Elaborate spectral reflectance models have been derived [9], [19], but they are all lacking real-time capacity or perhaps even general representation of spectral reflectance properties. Besides increasing realism it is possible to scale the amount of information contained in an image. Additional details may gradually be revealed as different lights illuminate the scene. The use of spectra increases the dimensionality of the output image. The actual look of the visualization may be determined by choosing an appropriate light. To make this choice more manageable our proposed 'light-dial' widget restricts the illuminant to be a combination of predefined light spectra. The light dial also offers a very simple interaction and exploration metaphor once the materials and lights are set up by an expert. So, the method could be very good for educational applications (medical and otherwise), museum exhibits, etc.

A. Conclusions

In volume rendering in general, the focus usually is not photo-realistic rendering, and that has not been the case in this paper either. Nevertheless, the method is based on physical models rooted in spectral measurements. We have shown how to use concepts such as metamers and colour constancy in order to explore a volumetric data set. These effects are not possible by using traditional RGB.

Based on an improved linear model, a practical framework for spectral volume rendering has been developed. To emphasize the use of re-illumination, we propose a method for rendering spectral scenes independent from a light source. This allows the user to obtain different impressions of the data under different lights, an enhancement of any tech-

nique that utilizes a transfer function for volume exploration. The spectral design method allows the preparation of specific illumination situations to highlight certain aspects of the data.

We have shown that our techniques are effective enhancements for transfer functions. While some of the demonstrated effects could perhaps be achieved by changing of opacities and assigned colours in the traditional RGB-based rendering system, a complete re-rendering of the scene becomes necessary. Post-illumination proves useful for the interactive exploration of visualizations, especially for high-quality rendering methods such as ray-tracing.

Acknowledgments: The *hipiph* dataset is courtesy of Scripps Clinic, California. The *uncbrain* is courtesy of UNC Chapel Hill. The *tomato* is from Berkeley Laboratory. The *frog* is part of VTK by Kitware (TM). The *engine* is available at www.volvis.org

REFERENCES

- [1] Steven Bergner, Torsten Möller, and Mark Drew. A spectral engine for visualization. Technical Report SFU-CMPT-TR2004-06, School Of Computing Science, Simon Fraser University, June 2004.
- [2] Steven Bergner, Torsten Möller, Mark S. Drew, and Graham D. Finlayson. Interactive spectral volume rendering. In *IEEE Visualization 2002*, pages 101–108. IEEE Press, October 2002.
- [3] Stuart K. Card and Thomas P. Moran. User technology: From pointing to pondering. In R.M. Baecker, J. Grudin, W.A.S. Buxton, and S. Greenberg, editors, *Readings in Human Computer Interaction: Toward the Year 2000*, pages 587–602. Morgan Kaufmann Publishers, San Francisco, California, 1995.
- [4] Robert A. Drebin, Loren Carpenter, and Pat Hanrahan. Volume rendering. In *Computer Graphics (Proceedings of SIGGRAPH 88)*, volume 22, pages 65–74, Atlanta, Georgia, August 1988.
- [5] Mark S. Drew and Graham D. Finlayson. Multispectral processing without spectra. *J. Opt. Soc. Am. A*, 20:1181–1193, 2003. Also: Representation of colour in a colour display system, UK Patent Application No. 0206916.9. Under review, British Patent Office.
- [6] K. Engel, M. Kraus, and Th. Ertl. High-quality pre-integrated volume rendering using hardware-accelerated pixel shading. In *2001 SIGGRAPH / Eurographics Workshop on Graphics Hardware*, pages 9–16. ACM Press, August 2001.
- [7] Alireza Entezari, Randy Scoggins, Torsten Möller, and Raghu Machiraju. Shading for Fourier volume rendering. In *Proc. of IEEE Volume Visualization and Graphics Symposium 2002*, pages 131–138, October 2002.
- [8] G.D. Finlayson, M.S. Drew, and B.V. Funt. Spectral sharpening: sensor transformations for improved color constancy. *J. Opt. Soc. Am. A*, 11(5):1553–1563, May 1994.
- [9] J.S. Gondek, G.W. Meyer, and J.G. Newman. Wavelength dependent reflectance functions. In *Proc. SIGGRAPH 1994*, pages 213–220, 1994.
- [10] Kazufumi Kaneda, Yoshinori Dobashi, Kazunori Yamamoto, and Hideo Yamashita. Fast volume rendering with adjustable color maps. In *1996 Symposium on Volume Visualization*, pages 7–14. ACM SIGGRAPH, 1996.
- [11] Steven Kiltath and Torsten Möller. Splatting optimizations. Technical Report SFU-CMPT-04/01-TR2001-02, School Of Computing Science, Simon Fraser University, April 2001.
- [12] J. Kniss, S. Premoze, Ch. Hansen, P. Shirley, and A. McPherson. A model for volume lighting and modeling. *IEEE Trans. on Visualization and Computer Graphics*, 9(2):150–162, April-June 2003.
- [13] Robert L. Mack and Jakob Nielsen. Usability inspection methods: Executive summary. In R.M. Baecker, J. Grudin, W.A.S. Buxton, and S. Greenberg, editors, *Readings in Human Computer Interaction: Toward the Year 2000*, pages 170–181. Morgan Kaufmann Publishers, San Francisco, California, 1995.
- [14] Tom Malzbender. Fourier volume rendering. *ACM Transactions on Graphics*, 12(3):233–250, July 1993.
- [15] Nelson Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, June 1995.
- [16] Herke Jan Noordmans, Hans T.M. van der Voort, and Arnold W.M. Smeulders. Spectral volume rendering. In *IEEE Transactions on Visualization and Computer Graphics*, pages 196–207. IEEE Computer Society, 2000.
- [17] Jeff Orchard and Torsten Möller. Accelerated splatting using a 3D adjacency data structure. In *Graphics Interface 2001*, pages 191–200, June 2001.
- [18] M.S. Peercy. Linear color representations for full spectral rendering. In *Computer Graphics (SIGGRAPH '93)*, pages 191–198, 1993.
- [19] Jos Stam. Diffraction shaders. In *Proceedings of SIGGRAPH 99*, Computer Graphics Proceedings, Annual Conference Series, pages 101–110, Los Angeles, California, August 1999. ACM SIGGRAPH / Addison Wesley Longman.
- [20] Y. Sun, F.D. Fracchia, M.S. Drew, and T.W. Calvert. A spectrally-based framework for realistic image synthesis. *The Visual Computer: International Journal of Computer Graphics*, 17(7):429–444, 2001.
- [21] Takashi Totsuka and Marc Levoy. Frequency domain volume rendering. In *Proceedings of SIGGRAPH 93*, Computer Graphics Proceedings, Annual Conference Series, pages 271–278. ACM SIGGRAPH / ACM Press, August 1993.
- [22] Lee Westover. Footprint evaluation for volume rendering. *Computer Graphics*, 24(4):367–376, August 1990.
- [23] G. Wyszecki and W.S. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulas*. Wiley, New York, 2nd edition, 1982.