

An Analytic Model for Full Spectral Sky-Dome Radiance

Lukas Hosek*
Charles University in Prague

Alexander Wilkie†
Charles University in Prague

Abstract

We present a physically-based analytical model of the daytime sky. Based on the results of a first-principles brute force simulation of radiative transfer in the atmosphere, we use the same general approach of fitting basis function coefficients to radiance data as the Perez and Preetham models do. However, we make several modifications to this process, which together significantly improve the rendition of sunsets and high atmospheric turbidity setups – known weak points of the Preetham model. Additionally, our model accounts for ground albedo, and handles each spectral component independently. The latter property makes it easily extensible to the near ultraviolet range of the spectrum, so that the daylight appearance of surfaces that include optical brighteners can be properly predicted. Due to its similar mathematical properties, the new model can be used as a drop-in replacement of the Preetham model.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture;

Keywords: skylight models

1 Introduction

Renderings of outdoor scenes usually include vistas of the sky. Since everybody is intuitively familiar with its appearance, a correct skylight model can often make the difference between a good and an uncanny looking image. More importantly, though, a radiometrically plausible model that provides reliable illuminance data is important for predictive rendering purposes in appearance-sensitive areas like architecture and illumination engineering.

One common approach to rendering scenes that are lit by a sky-dome luminaire is to use skyboxes, or some other form of environment map with actual photographs of the sky, as light sources. This approach produces visually pleasing results, but has some disadvantages: you cannot easily alter the used HDR photo of the sky when you decide the sun needs to be just a little bit lower and the sky needs to be just a little bit hazier.

The other approach is to use an analytical model, ideally one that can be fully tweaked to suit the needs of the artist. The intention of this work is to produce such an analytic model. Our work is based on the widely used Preetham model, but it fixes some of its most apparent flaws, allows for the simulation of a wider variety of atmospheric conditions, and offers additional features.

In order to be able to develop such a model, we first had to obtain reference data. In the case of skylight, capturing the large amounts of data needed for this from nature in a reliable and repeatable fashion is very hard, which is why we resorted to implementing a brute force, first-principles model of atmospheric light transport. This model, which we describe in section 3, was used for generating a large number of reference images of the sky-dome. Then, using this reference data, and inspired by how the Perez and Preetham models are derived, we devised an extended skylight model, and fitted the parameters of the formulas to our reference renderings through an

offline optimisation process. The results of this fitting process are discussed in section 5, and a reference implementation that includes the numerical data we obtained is included as supplemental material to this paper. Since our proposed model fixes most known flaws of the widely used Preetham model without being significantly more complex, we hope that our findings are useful to all those who require a flexible, realistic model of sky-dome luminance.

2 State of the Art in Sky Models

2.1 Luminance-only Sky Models

Especially for earlier applications in architecture and lighting design, simpler models that only describe the luminance distribution of a clear sky were devised, and even standardised by the CIE. They can not be directly used for renderings, but are to this day still very useful for their original purpose. They also served as starting points for some of the sky models that include colour information.

An analytic formula for describing the luminance of clear skies with low turbidity was devised by Perez et al. [1993]:

$$\mathbb{F}_{Perez}(\theta, \gamma) = (1 + Ae^{B/\cos\theta})(1 + Ce^{D\gamma} + E\cos^2\gamma) \quad (1)$$

This formula was later adapted by CIE in a slightly different form:

$$\mathbb{F}_{CIE2003}(\theta, \gamma) = (1 + Ae^{B/\cos\theta})(1 + C(e^{D\gamma} - e^{D\frac{\pi}{2}}) + E\cos^2\gamma) \quad (2)$$

In these formulas, γ is the angle formed by the view direction and a vector pointing towards the sun, and θ is the angle formed by zenith and view direction. The values returned by these formulas are normalised so that the value at zenith is one and then multiplied by the luminance at zenith obtained from elsewhere, which gives the final *Luminance distribution function*:

$$Y = \frac{\mathbb{F}(\theta, \gamma)}{\mathbb{F}(0, \theta_S)} \cdot Y_Z \quad (3)$$

where θ_S denotes the angle formed by Sun and zenith and Y_Z is the luminance at the zenith.

The parameters A through E in the Perez formula and its modified CIE counterpart are used to tune the luminance distribution and do not directly translate to any physical quantities. One way to obtain these parameters is by using tabulated values. This was done in the joint ISO/CIE standard [ISO/CIE 2004], which comprises 16 different models describing luminance distribution for different atmospheric conditions that vary from clear sky to overcast with different luminance turbidities. The CIE model is targeted mostly at architectural applications and does not provide chroma data.

2.2 Sky Models that include Colour Data

One of the first coloured skylight models that is directly useful for rendering purposes was proposed by Nishita et al. [1993]. It simulated the appearance of the sky due to single scattering only, and ignored inter-reflections between the ground and air molecules. A revised model was presented by Nishita et al. [1996]; it used a discretised pre-computation to simulate multiple scattering.

*e-mail: hosek@cgg.mff.cuni.cz

†e-mail: wilkie@cgg.mff.cuni.cz

A more recent physically based model was presented by Haber et al. [2005]. The authors follow the same basic principle used in Nishita’s paper – subdivision of the atmospheric body into discrete blocks and calculation of radiative transfer between them. Given the correct configuration of atmospheric constituents, it can realistically predict appearance of the sky over a wide range of atmospheric conditions. In the form it was presented, the model does not account for the influence of ground reflectance, although according to the authors, it could easily be modified to do so.

A model that incorporates multiple scattering, including scattering in water bodies, and that allows for real-time rendering by utilising a set of lookup tables was presented by Elek and Kmoch [2010]. This model uses subdivision of the atmospheric body into discrete blocks and works in two phases – the first, time costly one, pre-computes a 4-dimensional look-up table which can then be used in the second phase in real time to render images from any position and at any incident light angle. This model is fully spectral and takes account of ground albedo. The value of ground albedo has to be decided in advance, since it is baked into the look-up table.

The widely used Preetham model [Preetham et al. 1999] is directly based on the model of Perez et al. [1993], but unlike in the CIE models, the parameters A to E are not tabulated, but calculated analytically. Their approach was to generate reference images using Nishita’s model and then, by non-linear optimisation, fit the Perez formulas to the reference images. The results were a set of linear functions that take one parameter, *turbidity*, as an argument, and return the parameters A to E , as well as a bi-cubic function of turbidity and solar elevation angle for calculating zenith luminance. In addition to luminance, the Preetham model also provides two chroma channels that are calculated using the same approach, and the authors also provide a way to convert the outputs to spectral radiance data. The resulting model is fast and easy to implement, and it is currently the de facto standard analytic model of spectral sky-dome radiance. Habel et al. [2008] reformulated the Preetham model using spherical harmonics, which can be used to efficiently compute the illumination of objects by the sky.

2.3 Describing the Atmosphere via Turbidity

Linke’s turbidity factor [McCartney 1976], or turbidity for short, serves as a simple and intuitive measure of the aerosol content of the air. It is defined as the ratio of the *additional* optical thickness of the atmosphere in question t_h to the optical thickness of an idealised atmosphere that consists only of pure gas t_m :

$$T = \frac{t_m + t_h}{t_m} \quad (4)$$

This allows the user to easily define sky appearance without worrying too much about the intricacies of meteorology: $T = 2$ yields a very clear, Arctic-like sky, $T = 3$ a clear sky in a temperate climate, $T = 6$ a sky on a warm, moist day, $T = 10$ a slightly hazy day, and values of T above 50 represent dense fog. Of course, flattening the diversity of all possible atmospheric conditions into one number requires one to make somewhat arbitrary decisions with regards to particle size distributions, and the shape of the vertical falloff of aerosol particle density. But since overall aerosol content is the dominant factor in sky appearance, it still is a highly useful parameter, which we also opted to use in our model.

2.3.1 Turbidity in the Preetham Model

For their model, the authors of the Preetham model decided to go with a valid turbidity range of 2 to 6. Since the Preetham model uses fitting to ad-hoc formulas, it can only work over a limited range of turbidities – and this particular range presents a trade-off between

usefulness and accuracy. Their fitting does work well for high solar elevation – low turbidity situations. However, partly due to the limitations of the Perez formula, and partly due to the simple linear functions chosen for calculating the parameters, it fails at the opposite part of the spectrum, and paints sunsets at high turbidities a uniform orange-yellow: a well-known deficiency of the Preetham model that can be seen in figure 1. A review of the Preetham model by Zotti et al. [2007] also shows other issues, namely that under specific conditions the model yields negative luminance values.

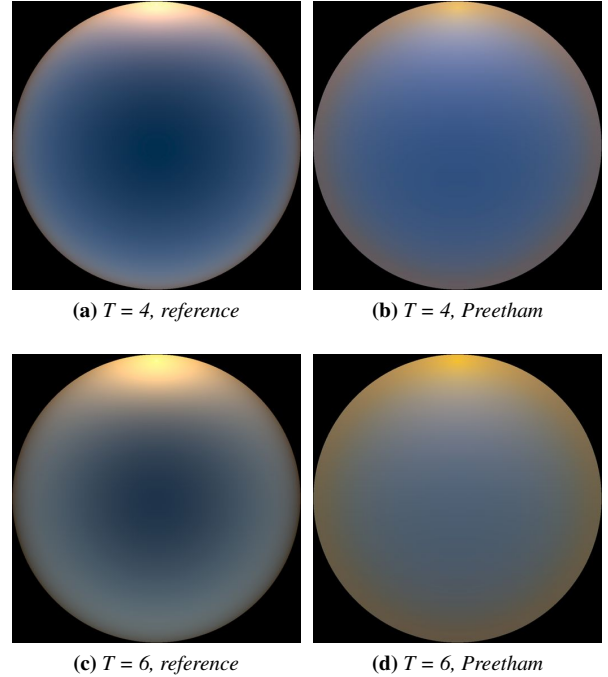


Figure 1: Comparison of reference solutions (computed according to the techniques outlined in section 3) to the Preetham model for turbidities 4 and 6 at a solar elevation of 4° . While the match is still reasonably good for $T = 4$, it already has visible deficiencies: the area around the sun is not bright enough, and the zenith is not dark enough. At $T = 6$, both the horizon colour pattern and the luminance distribution are considerably off.

3 Sky Dome Reference Data Generation

As already mentioned in the introduction, capturing sky dome radiance data for a representatively large number of atmospheric conditions and solar elevations from nature in a reliable, controlled and repeatable fashion is a hard problem. So we opted to obtain this data via a brute force, first-principles model simulation of atmospheric light transport instead. The underlying algorithm of our simulation is a brute force path tracer that is capable of simulating interaction between light and participating media. In the following sections we give a brief overview of the techniques used by this path tracer.

3.1 Atmospheric Scattering

The propagation of sunlight through a cloudless atmosphere is determined by scattering events in the atmosphere itself, and interactions of the light with the ground. Interactions with the ground are simple to model, and can be described by standard models of reflectance. For the purposes of light propagation, the atmosphere itself can be thought of as being made up of two constituents:

aerosols, and the air molecules themselves. The interaction of light with these is somewhat more involved, and covered by Lorenz-Mie and Rayleigh scattering theory, respectively.

3.1.1 Scattering Event Generation

The main technical issue in a path tracer used for atmospheric simulation is the efficient generation of scattering events. We use a technique for unbiased generation of free-flight paths in non-homogeneous environment as described in Raab et al. [2008].

3.1.2 Rayleigh Scattering

Rayleigh scattering describes the interaction between electromagnetic radiation and particles significantly smaller than the wavelength of the radiation – in our case visible light and the gas phase constituents of the atmosphere. Rayleigh scattering is elastic: meaning no energy is transferred between the interacting particles. Translated into path tracing terms, this means the energy of a ray does not change, although its direction can. Rayleigh scattering is the reason why the sky is blue, and the red tint of setting sun is also caused by Rayleigh scattering: blue and violet light is more likely to be out-scattered from the direct path between an observer and the setting sun. The effective cross-section of Rayleigh scattering is strongly wavelength dependent, about 16 times higher for the violet end of spectrum than for the red end, which means that along its journey through the atmosphere, violet and blue rays have a much higher chance of being scattered towards the observer.

The macroscopic cross-section of Rayleigh scattering can be described by

$$S_R^{macro} = \frac{128\pi^5 \alpha_0^2 (6 + 3\rho_n)}{3\lambda^4 (6 - 7\rho_n)} \quad (5)$$

where α_0 is the volume polarisability (e.g. $1.7 \cdot 10^{-24} \text{ cm}^3$ for terrestrial air), ρ_n is depolarisation factor (0.0279 for naturally polarised light) and λ is the wavelength. We assumed that the particle density decays exponentially with altitude, with with half-life of the exponential density $t_{half} = 7994\text{m}$.

Once a Rayleigh scattering event occurs, the ray is recast in a new direction, with the distribution of directions governed by the phase function:

$$F_R(\theta) = \frac{3}{4}(1 + \cos^2(\theta)) \quad (6)$$

with θ as the angle between outgoing and incoming light directions.

In a simulation with just Rayleigh scattering enabled, the sky colour outside the solar disc is restricted to shades of blue, with the hue corresponding to that of a real sky on a cloudless day. The blue hue is a combined effect of the Sun’s spectrum and the fourth-power dependence of the Rayleigh scattering on the wavelength. When the Sun is placed at altitudes lower than about 10° , the effect of the dependence of the free-flight path on wavelength becomes visible. At sea level, the free flight path of violet light is about 50km, in contrast to red light with a free flight path of about 200km. This means that red light is significantly less likely to be out-scattered, which produces the distinctive red tint of a setting sun, and also a slight reddening of the sky above the horizon. Our reference model reproduces all these effects. The directional dependence of intensity of scattered light in equation 6 suggests the sky will be darkest at 90° away from the sun, but due to multiple scattering, that effect is not as pronounced as the equation would suggest.

3.1.3 Lorenz-Mie Scattering

Lorenz-Mie theory gives us a tool to predict the contribution of particles of sizes comparable to the wavelength of visible light, most

notably water droplets. The interaction is again elastic. Its phase function is strongly anisotropic, and Mie scattering strongly favours light scattered in the forward direction.

Unlike Rayleigh scattering, where the phase function always has the same shape, the pattern of Mie scattering and its macroscopic cross-section varies with droplet size and wavelength. In our simulation, we used droplet size distributions by Eldridge [1961]. That paper also includes meteorological range measurements, which can be used to approximately calculate turbidity. The measurements were taken in a dense fog, corresponding to turbidity of about 60. For our simulation runs with different turbidities, the particle concentrations were scaled accordingly. A further difference to Rayleigh scattering is that the free-flight path of Mie scattering is not strongly wavelength dependent. As stated above, the shape of the phase function varies with wavelength, though, which for larger water droplets can manifest itself in rainbow effects that are caused by a slight shift in the position of the secondary peak. This is however not an issue for the particle sizes we simulate ($< 20\mu\text{m}$).

Unfortunately, Mie scattering theory does not come with a nice set of closed form formulas like Rayleigh scattering does; both effective cross-section and intensity can only be described in the form of a infinite expansion series [Hulst and Hulst 1957]. For our simulation we used pre-calculated tabulated values for cross-sections [Mätzler 2002]. Once again, we assume density decays exponentially with altitude, with $t_{half} = 1200\text{m}$. The complex shape of the phase function is modelled using a modified version of the Henyey-Greenstein analytic approximation of the general Mie phase function [Cornette and Shanks 1992]:

$$F_M(\theta) = \frac{3(1 - g^2)}{2(2 + g^2)} \cdot \frac{1 + \cos^2(\theta)}{(1 + g^2 - 2g \cdot \cos(\theta))^{\frac{3}{2}}} \quad (7)$$

where g is the anisotropy factor.

3.2 Path Tracer Implementation Details

We implemented the brute force path tracer used to generate the reference images in C++. The simulated world geometry consists of just two objects – a sphere that represents the Earth, with a Lambertian diffuse surface of a user-supplied colour, and another sphere that represents the upper boundary of the atmosphere, placed 100km above sea level. When the ray hits the upper boundary, it is considered to have exited the atmosphere and is tested for collision with the sun. The spectral radiance values of the sun were taken from Preetham et al. [1999]. The camera is placed 10m above sea level, and uses a circular fisheye projection that is oriented towards the zenith. For the purposes of generating visualisations for this paper, the spectral results were converted to CIE XYZ values and then tone-mapped using the “Interactive Calibration” algorithm proposed by Matkovic and Neumann [1996].

The sun is the only source of light in this setup. Since it covers only $\approx 0.0005\%$ of the total celestial sphere, in a primitive implementation an overwhelming majority of the rays would be wasted. One way to alleviate this problem is to render more images at once. Suppose we want to render images of the sky with the solar elevation ranging from 0° to 90° in 10° increments. We place 10 suns at their respective positions, and when a ray exits the atmosphere and hits any of the suns, information on which sun was hit is passed to the ray sampler along with the radiance information. The ray sampler then splats the sample only to the image associated with the sun that was hit. This approach yields a $10\times$ speed increase for free. This idea can be further extended by having multiple suns for the same elevation placed in circles around the zenith. Once a path that hits a sun is generated, it is rotated so that the sun moves to its canonical position and the pixel is then splatted onto the corresponding




Figure 2: Brute force path tracer renderings: a demonstration of the difference varying turbidity values make to the appearance of the sky at daybreak. Solar elevation is 4° . Note that these images only show the sky-dome without the solar disc itself.

place on the projection plane. When computing an image for solar elevation 0° , we can e.g. place 720 suns around the horizon without overlapping (since each of them is 0.5° wide), and achieve a speedup of up to $720\times$ if we use a suitable acceleration structure for intersection testing. We did not use direct sampling of the sun (i.e. light source sampling), since the simulator is fast enough as a reference solution generator with the described “many suns” optimisation. Also, in its current simple form the simulator makes only minimal assumptions about the atmospheric density distribution and the nature of the scattering events that can occur. This is a desirable feature for further research, in which we plan to investigate more complex atmospheric distributions and environments.

The brute force model provides us with reference images, but it was always clear that it would be totally unusable for any sort of production work – a sufficiently converged image requires shooting about one million rays per pixel and calculating a 128×128 image takes somewhere between 40 minutes and 3 hours, depending on the turbidity, on a 2.66GHz Core i7 920 machine. In addition to this, the simulation requires that one supplies it with physical parameters such as *particle polarisability* or *anisotropy factors*, which are not the sort of controls one would want to force on anyone who merely wants to use a sky-dome model. What is useful in practice are analytical models with just a few easily comprehensible and expressive parameters. All results from such a model have to be obtainable by using closed form formulas, which can be tuned by a reasonably small number of comprehensive parameters.

4 Reference Data: Simulation Results

The simulation described in the previous section produces images with the distribution patterns expected from a physically-based simulation: an example is shown in figure 2. We systematically per-

formed these simulations for a wide range of input values for wavelength, solar elevation, turbidity, and ground albedo. The two most interesting features of the result images are the effect of ground albedo on the radiance distribution, and the colour of the sunset.




Figure 3: Brute force path tracer renderings. Difference between low and high ground albedo α , solar elevation 40° , $T = 4$. Images rendered in linear colour space.

4.1 The Effect of Ground Albedo

Something which apparently has not been discussed much in graphics literature is the influence of ground albedo on the overall appearance of the sky. High albedo values can e.g. occur in winter scenes - snow reflectance reaches values up to 1 at 550nm, so almost all incident radiation is reflected and backscattered towards the viewer due to Mie scattering. On the other hand, grass has albedo of about 0.09 at 550nm. Especially in high turbidity settings, changing ground albedo alters the brightness of the whole sky-dome in a perceptible fashion. Figure 3 shows that a highly reflective surround can almost double sky dome brightness.

This observation was actually our initial motivation for developing a modified version of the Preetham skylight model. The other benefits of our new model, such as the more realistic sunset colours, or the better handling of high turbidity situations, emerged as further results once we started work on the modification process. But as figure 10 shows, the changes in sky-dome appearance due to ground albedo are not restricted to monochrome brightness changes.

4.2 Sunset Colour

The second important issue that we noticed in the brute force reference renderings was that the predictions of the Preetham model for sunset situations are actually not particularly good. Images 4c and 4d show a comparison between a fisheye photo of a real sunset sky, and our reference rendering. Figure 1 shows side by side comparisons of our analytical model (which is visually indistinguishable from the reference rendering) to the Preetham model for similar turbidity values, and solar elevations. Both the absolute colours, as well as the distribution of the colours in the Preetham model can only be considered a rough approximation at best.

5 An Improved Analytical Skylight Model

Based on the results discussed in the two previous sections, we then proceeded to derive an analytical model that is fitted to these results. Our goals were threefold: to obtain a genuinely spectral analytical model with separate fitting for different wavelengths, to have this fitting process start in the near ultraviolet range to enable proper outdoor fluorescence rendering, and to extend the useful range of turbidities to 10. We did not go beyond a turbidity of 10 mainly because the layering of very hazy atmospheres starts to be rather specific after some point: inversion layers, which would require additional parameters to control, can start to play a visually prominent

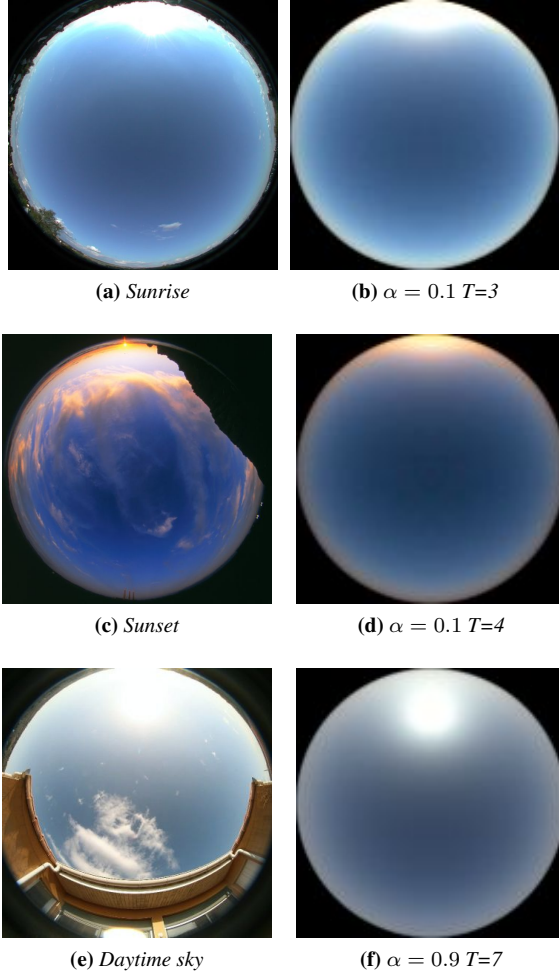


Figure 4: Comparison between fisheye photos (left) and brute force path tracer outputs (right). *4c* was taken on Velvia film, which produces characteristically saturated pictures. To compensate, the path traced image was rendered with a white point of 4400K.

role. We therefore restricted ourselves to a scenario of a gradual increase in haze up to moderate haze levels. A CIE XYZ fitting is also available for applications where a full spectral model is not needed.

Using the brute force model, we generated a reference dataset for turbidity ranging from 1 to 10, albedo 0 to 1 and eleven spectral channels in total: nine in the visible range (400nm - 720nm in 40nm increments) and two in the ultraviolet range (320nm and 360nm). For each of these combinations of turbidity, albedo and wavelength, 46 images with solar elevations ranging from 0° to 90° were generated. The appearance of the sky undergoes a very dramatic change as the sun goes through the first few degrees of sunrise. After the solar elevation reaches about 15°, sky-dome appearance stabilises and subsequent changes are minimal.

To better capture this behaviour, the reference images were not generated at equidistant solar elevations. Instead, more samples were taken at lower solar elevations: the n -th image ($n \in \{0, 1 \dots, 45\}$) was placed at solar elevation $(\frac{n}{45})^3 \cdot 90^\circ$. The method for calculating radiance distribution parameters, described further in this section, directly reflects this decision.

The path tracer was configured to exclude rays that directly hit the

sun without being scattered first, since this would produce a discontinuity in the data: the solar disc is several orders of magnitude brighter than the surrounding sky. As with the Perez model, the fitted model therefore does not include the solar disc, so users of the model have to include it separately.

5.1 Extending the Perez Model

The Perez formula can capture luminance distributions of skies with low turbidities almost flawlessly, providing the A, B, \dots, E parameters are fitted specifically to the setup we want to simulate (i.e. including a specific solar elevation). That is because at its core, the Perez formula is basically a single scattering model with Rayleigh phase function and exponential out-scattering. As the turbidity rises, the contribution of Mie scattering becomes more pronounced, though. Recall that Mie scattering is highly anisotropic.

The zero order glow of Mie scattering produces a phenomenon called *circumsolar ring* (aureole). This produces a highly localised spike, one that is impossible to accurately fit using the original Perez formula, which is also at the core of the Preetham model. Our first goal was to replicate this phenomenon and for that, we used a modified version of the Perez formula which includes an anisotropic term χ (the reader may find this term similar to the phase function of Mie scattering we used in the brute force path tracer) that places a localised glow around the solar point.

Another thing apparent from viewing reference images and fisheye photographs of the sky is that at lower solar elevations, the aureole does not extend towards the zenith nearly as much as it does towards the horizon and sides. There are obvious reasons for this: the higher the viewing angle, the fewer scatterers are in the way to in-scatter light in the viewer’s direction. We added a term $I \cdot \cos^{\frac{1}{2}} \theta$, which produces a smooth gradient around the zenith. When I is set to a negative value (as it always is with our fitted data), it suppresses brightness around the zenith, and thereby also reduces the extent of the aureole around the zenith.

Finally there is one more minor technical issue: in the original Perez formula, as the viewing direction neared the horizon, the values diverged towards infinity and when querying luminance exactly at the horizon, it would produce division by zero: when θ becomes 90°, $B / \cos \theta$ becomes undefined. This is solved simply by adding a small fudge factor that moves the offending division by zero about 0.5° below horizon and also solves the extremely bright rim around horizon. The final formula is:

$$\mathbb{F}(\theta, \gamma) = (1 + Ae^{\frac{B}{\cos \theta + 0.01}}) \cdot (C + De^{E\gamma} + F \cos^2 \gamma + G \cdot \chi(H, \gamma) + I \cdot \cos^{\frac{1}{2}} \theta) \quad (8)$$

$$\chi(g, \alpha) = \frac{1 + \cos^2 \alpha}{(1 + g^2 - 2g \cdot \cos \alpha)^{\frac{3}{2}}} \quad (9)$$

The inputs of this formula are the same as those in the original Perez formula: θ is the angle formed by viewing ray and zenith, and γ is the angle between viewing ray and the solar point.

Notice that compared to the original Perez formula in equation 1, the value of 1 at the beginning of the second parenthesis block has been replaced by a variable C . This is due to another complication that arose with adding an anisotropic term. Recall that in the original Perez model, values of the \mathbb{F} function were normalised against zenith luminance (see equation 3). Our model can now reproduce the huge spikes in brightness around the sun, which poses a problem with very high solar elevations ($> 85^\circ$), because the function




Figure 5: *The coordinate system used in the model.*

that calculates luminance (or spectral radiance, as would be the case in our model) at zenith would now also have to reproduce the spike.

Instead of devising functions that would replicate the spikes in luminance (or, in our case, spectral radiance), we opted out of any explicit normalisation in the model. The final spectral radiance is now calculated using

$$L_\lambda = \mathbb{F}(\theta, \gamma) \cdot L_{M\lambda} \quad (10)$$

where $L_{M\lambda}$ is the expected value of spectral radiance in a point randomly picked in the upper hemisphere with uniform distribution. This implies that the expected value of the \mathbb{F} function is 1 – but this is due to the fitting rather than due to an explicit scaling. The main idea stays the same – we compute the *distribution* \mathbb{F} and one *master value* separately, only this time there is no sudden, hard to replicate spike in the master value, making it much easier to fit. Since the distribution function is now not anchored by the zenith value, it has to be accurate in absolute terms (as opposed to the Preetham model, where luminance values have to be accurate relative to zenith). Hence the additional C variable in equation 8.

5.2 Exclusion of Dusk Conditions

Even though having a single, unified sky-dome model that also includes dusk conditions (e.g. solar elevations of -10°) would be useful, there are technical reasons why we did not extend the fitting process beneath solar elevations of 0° . Figure 6 shows a reference image for a solar elevation of -5° . For the fitting process to the function we just described (i.e. an improved Perez model), the additional dark crescent-shaped region in the lower part of the image would pose a substantial problem. One would in all likelihood have to extend the function described in the previous section yet again to capture this additional feature in the luminance distribution, which would in turn substantially complicate the fitting process.




Figure 6: *An over-exposed reference rendering of the sky at solar elevation -5° . The crescent-shaped shadow which the Earth starts to cast into the atmosphere once the sun has gone beneath the horizon can be seen in the lower part of the image.*

5.3 Separate Fitting for Reference Data Channels

The new model comes in two versions: one fully spectral, and one that supplies CIE XYZ values. In both versions, the model produces radiance data that is separately fitted for each channel provided in the reference data, be that spectral values, or CIE XYZ colours. The rationale behind the separate fitting for each band is obvious from figure 7: no single set of coefficients could properly capture the variations in pattern seen across the spectral range.




Figure 7: *Brute force simulation results for the range from UV to the red part of the spectrum for three solar elevations. Note the substantial changes in the patterns, in particular for the UV range. Fitting all these patterns to a single set of coefficients with a function like the one used in the Perez model (which was originally only intended for a single luminance channel) is bound to fail.*

The separate fitting for each wavelength band or colour channel of course means that when using the model, one has to calculate formula 10 separately for each band/channel, and supply its sub-functions \mathbb{F} and $L_{M\lambda}$ with different parameters (which we provide) for each band. This is the price one has to pay for accurate luminance data and chroma, (in the spectral version) the UV capability needed for handling of optical brighteners (UV is just another spectral band, after all), and most importantly, the ability to control ground albedo separately for each spectral band / colour channel.

5.4 Radiance Distribution Parameters

Radiance distribution parameters are the variables A, B, \dots, I that plug into equation 8. In the Preetham model, the parameters were calculated as a linear function of turbidity. We propose a different approach that uses Bezier curves for calculating these parameters. Again, figure 7 is instructive in motivating this: the luminance distributions for the various wavebands not only look appreciably different; they also *change in a different, non-linear fashion* as the solar elevation is altered. Having distribution parameters change with solar elevation is a major part contributing to the enhanced fidelity of the new model.

For each given wavelength, the parameters are calculated as follows: a 4-dimensional ($10 \times 2 \times 9 \times 6$) table M^λ of values $m_{T,\alpha,p,c}$ is provided – for each of the ten integer turbidities $T \in \{1, \dots, 10\}$, two albedo values $\alpha \in \{0, 1\}$ and nine distribution parameters $p \in \{1, \dots, 9\}$ there is a set of six control points $c \in \{1, \dots, 6\}$. These points are used for interpolation along a quintic bezier polynomial, with solar elevation η used as the interpolation parameter.

First, the solar elevation parameter is transformed into a $[0, 1]$ interval: $x = \sqrt[3]{\eta/\frac{\pi}{2}}$. As discussed before, most of the changes in the radiance distribution pattern happen abruptly at low solar elevations, therefore the cube root is used to spread them more evenly along the whole interpolation interval. The resulting vector of distribution parameters is calculated as follows:

$$\begin{aligned}
 v_p^\lambda = & m_{T,\alpha,p,1}^\lambda \cdot (1-x)^5 + \\
 & m_{T,\alpha,p,2}^\lambda \cdot 5x(1-x)^4 + \\
 & m_{T,\alpha,p,3}^\lambda \cdot 10x^2(1-x)^3 + \\
 & m_{T,\alpha,p,4}^\lambda \cdot 10x^3(1-x)^2 + \\
 & m_{T,\alpha,p,5}^\lambda \cdot 5x^4(1-x)^1 + \\
 & m_{T,\alpha,p,6}^\lambda \cdot x^5
 \end{aligned} \tag{11}$$

The result is a 1-dimensional, 9 element parameter vector V^λ – the wavelength, turbidity, albedo and elevation-specific vector of parameters to be used in formula 8. The arrays were fitted to the outputs of the reference path tracer using Levenberg-Marquardt non-linear least-square fitting in MATLAB.

If values for non-integer intermediate T and α are required, they can be obtained by calculating parameter vectors for adjacent turbidities and albedos and then linearly interpolating. Using interpolated parameter vectors generally produces results indistinguishable from generating full images from adjacent parameter vectors and then interpolating them. One potential problem would be that the optimisation process could choose to use completely different strategies for representing adjacent turbidities or albedos, in which case the interpolated results would be degenerated. To ensure this doesn't happen, interpolated results were included in the error function during the fitting process.

5.5 $L_{M\lambda}$ calculation

We use the same approach as with the radiance distribution parameters, only this time the result of the Bezier interpolation is just one number. A 3-dimensional ($10 \times 2 \times 6$) table R^λ of values $r_{T,\alpha,c}$ is provided and the same quintic bezier interpolation is used to get a single number $L_{M\lambda}$ – the expected value of radiance that is plugged into formula 10. The values in the table were again fitted using Levenberg-Marquardt and intermediate values can be obtained simply by interpolating.

6 Error Analysis

Due to the high dynamic range nature of the data generated by the reference path tracer, we use the logarithmic signal-to-noise ratio (SNR) calculated using the $20 \cdot \log_{10}$ rule as the measure of how well the new model approximates our reference data. Larger values of SNR correspond to better reproduction, and a SNR increase of 6db halves the average deviation. In the photographic community, a SNR of 20db is considered acceptable quality. Our path traced reference results are inherently somewhat noisy, so the SNR values quoted in this section tend to under-estimate the actual result quality a bit. Complete comparison charts for each spectral band with SNR data for each elevation are included as supplementary material.

6.1 Comparison with the Original Perez Function

Apart from analysing the performance of the new model, we also fitted the unmodified Perez formula to our reference data. This was done in order to validate the changes we made to the Perez sky dome radiance function that are discussed in section 5.1. It should be noted that the SNR results we show in table 1 are for the new model “as is” (i.e. based on a limited set of coefficient data, and with the interpolations contained in its design), while the parameters for the Perez formula used in this comparison were fitted

| | New model | | | | Perez model | | | |
|-------|-----------|------|----|--------|-------------|------|----|--------|
| | avg | min | T | η | avg | min | T | η |
| 320nm | 25.5 | 18.7 | 10 | 0°12' | 17.4 | 15.3 | 5 | 0° |
| 360nm | 26.2 | 19.8 | 6 | 0°1' | 16.1 | 14.7 | 2 | 1°42' |
| 400nm | 25.5 | 20.4 | 3 | 0°1' | 16.9 | 13.4 | 4 | 2°42' |
| 440nm | 24.1 | 18.4 | 10 | 7°54' | 17.4 | 12.4 | 8 | 0° |
| 480nm | 22.5 | 15.2 | 10 | 7°54' | 17.5 | 12.8 | 8 | 0°12' |
| 520nm | 21.0 | 13.0 | 10 | 7°54' | 17.4 | 12.2 | 9 | 5°45' |
| 560nm | 19.5 | 11.3 | 10 | 5°45' | 17.4 | 10.4 | 9 | 5°45' |
| 600nm | 18.3 | 10.0 | 10 | 5°45' | 17.0 | 8.5 | 10 | 4°2' |
| 640nm | 17.2 | 9.1 | 10 | 5°45' | 16.3 | 7.7 | 10 | 4°2' |
| 680nm | 16.3 | 8.2 | 10 | 4°2' | 16.1 | 7.0 | 10 | 4°2' |
| 720nm | 15.4 | 7.6 | 10 | 4°2' | 15.8 | 6.7 | 10 | 4°2' |
| CIE X | 20.0 | 11.2 | 10 | 5°45' | 18.0 | 11.1 | 10 | 5°45' |
| CIE Y | 20.0 | 11.5 | 10 | 5°45' | 17.9 | 11.5 | 10 | 5°45' |
| CIE Z | 24.0 | 17.1 | 10 | 7°54' | 18.1 | 17.0 | 10 | 7°54' |

Table 1: For each waveband/colour channel, this table lists the average SNR in db, as well as the worst SNR together with the turbidity and solar elevation where the worst case occurs. Also included are the average and worst case results for the per-waveband fittings done with the original Perez formula. The only waveband for which the original formula slightly outperforms the modified version used in the new model is 720nm, but the advantage is not significant. For all others, it is clearly inferior. The fittings with the worst SNR are shown in figure 8.

specifically for each waveband/turbidity/albedo/solar angle reference dataset. Because of this, the comparison is actually pretty unfair to the new model, but the original Perez formula still consistently performs worse than the modified formula, except for the average of the 720nm waveband, where it is slightly better.

All the worst cases are for albedo 0. At albedo 1, the model performs slightly, but without exception, better. The worst offenders for each spectral band are shown in figure 8. Generally, the performance gets worse with higher turbidities – with the exception of bands 360nm and 400nm, where the error was more affected by actual noisiness of the reference image than an inability to reproduce the shape correctly.

6.2 Colour Renderings, Worst Case Scenario

For this test, spectral model outputs for each turbidity/albedo/solar elevation configuration were rendered, then converted to CIE XYZ via the CIE standard observer colour matching functions and then converted from XYZ to L*a*b*. Results were compared to reference images, which were processed the same way and the signal to noise ratio between them was calculated. The worst offender is an image rendered at $T = 10, \alpha = 0, \eta = 5°45'$ with SNR = 12.2dB. The same test was also done for the XYZ fitting, where the worst performer was the image rendered at $T = 10, \alpha = 0, \eta = 4°51'$ with SNR = 10.8dB. Both of the images converted to sRGB are shown in figure 8.

7 Results

We performed the described fitting process for turbidities 1 (ideal gas atmosphere with no water droplets) to 10 (light haze) and the result provides a good match for the whole range. The new model produces more realistic localised aureoles and sunrise/sunset colours. The proposed method of calculating radiance distribution parameters proves to be very flexible, matches the reference data excellently and is usable for a much wider spectrum of conditions than the Preetham model. We found that the model can recreate any atmospheric condition up to moderate haze sufficiently well. Fur-





Figure 8: The lowest SNR fittings of the new model for each waveband of the spectral model, as well as for both the overall spectral and CIE XYZ versions of the model. **Top:** reference results. **Bottom:** output of the new model. These images correspond to the lowest SNR cases mentioned in table 1. Note that for very short wavelengths, the noise present in the reference data is a major factor in low SNR readings.

thermore, the new model can be directly used to provide UV radiance data without any modifications. This ability is potentially important for situations where the day-time appearance of objects that contain optical brighteners has to be computed. Figure 11 demonstrates the impact that inclusion of UV radiation can have on outdoor renderings of objects that use fluorescent colours.

As supplementary material to this paper, we provide an ANSI C reference implementation of our model which includes the matrices M^λ for both spectral and CIE XYZ data. A detailed comparison between the new model outputs (using the interpolation process for calculating distribution parameters), the original Perez model (with distribution parameters fitted specifically for each configuration) and reference path tracer outputs can be found in the supplementary material.



(a) Reference photograph

(b) Our model, $T=9$

Figure 9: A hazy sunrise. Note that the new model includes the dark rim that appears around the horizon at higher turbidities.

For programmers with a working implementation of the Preetham model at hand, the modifications we propose require only a comparatively small amount of work. The \mathbb{F} formula is slightly changed, but its principle of operation remains the same. The only completely new thing is the Bézier interpolation used for calculating mean radiances and distribution parameters.

Of all the modifications to the model, the $\cos^{\frac{1}{2}}$ term placed around the zenith turned out to be the biggest improvement. It alone can cut the minimum attainable mean square error during the fitting process by half. The model now accounts for ground albedo and can model scenes with arbitrary average spectral reflectance of the surrounding countryside. The influence of albedo on the appearance of the sky seems to have gone mostly unnoticed in the graphics community until now, even though everyone who finds himself e.g. in the desert tends to notice the peculiar milky-blue colour of the sky typical for such regions. Figure 10 shows the subtle but noticeable effects of ground albedo on the sky over a reconstructed archaeological site.

8 Conclusion and Future Work

In this paper, we have presented an analytic model that is based on the widely used Preetham model, but that offers a substantially higher degree of realism in three important points:

1. Due to the separate fitting process for each wavelength, and for each channel of CIE XYZ in the colour-space version of the model, the colour distribution on the sky-dome is much more realistic.
2. The luminance distribution is much more realistic for conditions of high turbidity.
3. The spectral version of the model also provides data for the near UV range, in order to handle daylight fluorescence.

The performance of the model was verified against a brute force simulation of atmospheric light transport. Due to the separate fitting for each waveband, our new model is slightly more costly to evaluate than the Preetham model. However, the difference is not large, and could be alleviated e.g. by pre-computation of sky maps as backdrops for the actual rendering process.

Future work will include an investigation of after-sunset conditions. As of now, we can recreate these conditions in the brute force simulation, but it cannot be fitted using the analytic model in its current form, because it cannot recreate the earth casting a shadow onto the atmosphere. Future work will also include an investigation of the polarisation patterns found on the sky-dome for conditions of low turbidity. So far, these patterns have only been handled in graphics on a rather ad-hoc basis [Wilkie et al. 2004], and our brute force simulator offers the possibility of explicitly fitting functions to the observed patterns. We also plan to investigate the effect that non-monotonic aerosol distributions in the atmosphere, which can e.g. be observed during an inversion weather pattern, have on the luminance patterns seen on the sky-dome.

References

- CORNETTE, W. M., AND SHANKS, J. G. 1992. Physically reasonable analytic expression for the single-scattering phase function. *Appl. Opt.* 31, 16 (Jun), 3152–3160.
- ELDRIDGE, R. G. 1961. A few fog drop-size distributions. *Journal of the Atmospheric Sciences* 18, 5, 671–676.
- ELEK, O., AND KMOCH, P. 2010. Real-time spectral scattering in large-scale natural participating media. In *Proceedings of Spring Conference on Computer Graphics 2010*, Comenius University, Bratislava, H. Hauser and R. Klein, Eds., 83–90.
- HABEL, R., MUSTATA, B., AND WIMMER, M. 2008. Efficient spherical harmonics lighting with the Preetham skylight model. In *Eurographics 2008 - Short Papers*, Eurographics Association, K. Mania and E. Reinhard, Eds., 119–122.

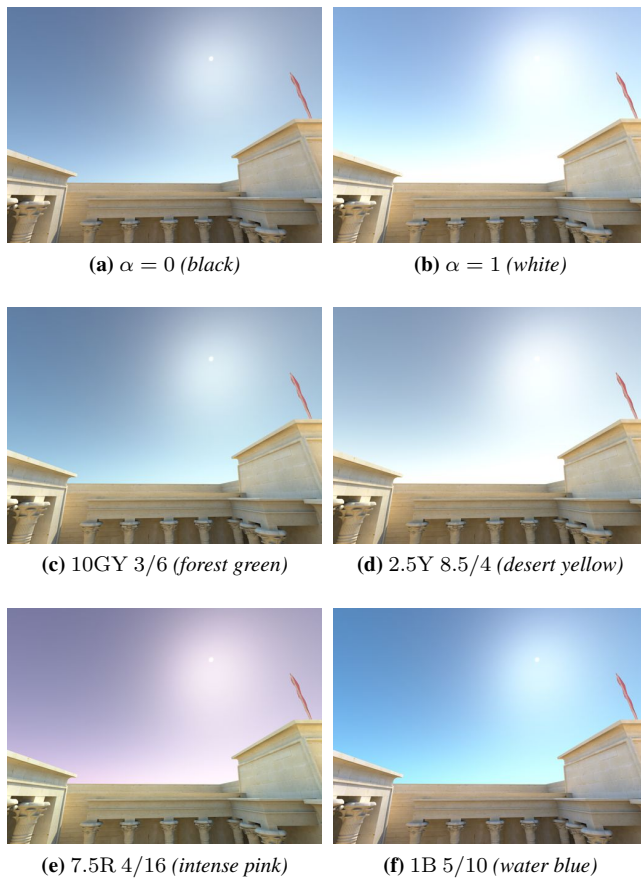


Figure 10: The influence of ground albedo on sky appearance for turbidity 5. The world outside the temple walls is assumed to be either achromatic with the given albedo, or of the specified Munsell colour. Reflectance data measured from a glossy Munsell Book of Color. Temple of Kalabsha model courtesy of Veronica Sundstedt.

HABER, J., MAGNOR, M., AND SEIDEL, H.-P. 2005. Physically-based simulation of twilight phenomena. *ACM Transactions on Graphics* 24 (October), 1353–1373.

HULST, H., AND HULST, H. 1957. *Light scattering by small particles*. Structure of matter series. Dover Publications.

ISO/CIE. 2004. Technical standard ISO 15469:2004 (CIE S 011/E:2003) "Spatial distribution of daylight – CIE standard general sky". Jointly published by ISO and CIE.

MATKOVIC, K., AND NEUMANN, L. 1996. Interactive calibration of the mapping of global illumination values to display devices. In *12th Spring Conference on Computer Graphics*, Comenius University, Bratislava, Slovakia, 155–161.

MÄTZLER, C. 2002. MATLAB functions for Mie scattering and absorption. *Research report 2002-8*, Institut für Angewandte Physik, Universität Bern, Switzerland.

MCCARTNEY, E. 1976. *Optics of the Atmosphere*. Wiley.

NISHITA, T., SIRAI, T., TADAMURA, K., AND NAKAMAE, E. 1993. Display of the earth taking into account atmospheric scattering. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, SIGGRAPH '93, 175–182.



Figure 11: The effect of including UV radiation in outdoor renderings that feature fluorescent surfaces. The inner objects have diffuse fluorescent BRDFs with measured re-radiation data of real materials. On the outer objects, only the non-fluorescent main diagonal of the same measurements is used. A significant number of fluorescent surfaces re-radiate energy from both the short-wavelength visible and the near UV ranges to longer wavelengths, and omitting UV irradiance alters their appearance. Illumination is with the proposed skylight model, turbidity 3. **Left:** 8 spectral samples, without UV **Middle:** 16 samples, with UV - notice how the leftmost two fluorescent surfaces gain in chroma **Right:** ΔE difference image

NISHITA, T., DOBASHI, Y., AND NAKAMAE, E. 1996. Display of clouds taking into account multiple anisotropic scattering and sky light. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, SIGGRAPH '96, 379–386.

PEREZ, R., SEALS, R., AND MICHALSKY, J. 1993. All-weather model for sky luminance distribution—preliminary configuration and validation. *Solar Energy* 50, 3, 235 – 245.

PREETHAM, A. J., SHIRLEY, P., AND SMITS, B. 1999. A practical analytic model for daylight. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, SIGGRAPH '99, 91–100.

RAAB, M., SEIBERT, D., AND KELLER, A. 2008. Unbiased global illumination with participating media. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*, A. Keller, S. Heinrich, and H. Niederreiter, Eds. Springer Berlin Heidelberg, 591–605.

WILKIE, A., ULBRICHT, C., TOBLER, R. F., ZOTTI, G., AND PURGATHOFER, W. 2004. An analytical model for skylight polarization. In *Rendering Techniques*, Eurographics Association, A. Keller and H. W. Jensen, Eds., 387–398.

ZOTTI, G., WILKIE, A., AND PURGATHOFER, W. 2007. A critical review of the Preetham skylight model. In *WSCG '2007 Short Communications Proceedings I*, University of West Bohemia, V. Skala, Ed., 23–30.