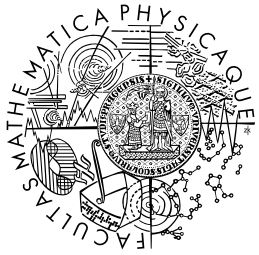

FACULTY
OF MATHEMATICS
AND PHYSICS
Charles University

DOCTORAL THESIS

Mgr. Lukáš Hošek

Atmospheric Rendering

Department of Software and Computer Science Education

Supervisor of the doctoral thesis:	doc. Dr. Alexander Wilkie
Study programme:	Informatics
Study branch:	Computer Graphics and Image Analysis

Prague 2019

I declare that I carried out this doctoral thesis independently, and only with the cited sources, literature and other professional sources.

I understand that my work relates to the rights and obligations under the Act No. 121/2000 Sb., the Copyright Act, as amended, in particular the fact that the Charles University has the right to conclude a license agreement on the use of this work as a school work pursuant to Section 60 subsection 1 of the Copyright Act.

In Prague, July 25, 2019 signature of the author

I would like to express my most sincere gratitude to my advisor, Alex Wilkie, for all his support and patience, and also for introducing me to the coolest hobby ever.

I would also like to thank Tom Bashford-Rodgers for helping with the atmospheric perspective fitting and Petr Vévoda for the path tracing optimizations.

Last but not least, I would like to thank my parents, Jana and Josef, who have always supported and encouraged me throughout my studies .

Title: Atmospheric Rendering

Author: Mgr. Lukáš Hošek

Department: Department of Software and Computer Science Education

Supervisor: doc. Dr. Alexander Wilkie, Department of Software and Computer Science Education

Abstract: The sky is an important feature of all outdoor scenes. This thesis explores the topic of skydome models — an approach to getting the sky and atmospheric effects integrated into a renderer. We discuss the physics of atmospheric scattering and describe in detail the construction of a first-principles path tracer atmospheric simulator implementation. The presented path tracer is fully spectral and produces polarization data. Finally, we present two different analytic skydome models. These analytic models are highly practical because they can be integrated even into real-time renderers and provide an excellent combination of fidelity and low computational cost. The first model is a simpler version, providing just the skydome spectral radiance. The second model also provides polarization data, after-sunset skies with an accurate modeling of Earth's shadow, aerial perspective and full sphere data.

Keywords: computer graphics, rendering, skylight models, atmospheric modeling

Contents

1 Introduction **3**
1.1 Our approach 4

2 Light transport **6**
2.1 The rendering equation 6
2.2 Participating media 7
2.2.1 Rayleigh scattering 7
2.2.2 Lorenz-Mie scattering 10
2.2.3 Atmospheric composition 14
2.3 Extending the rendering equation to account for participating media 15

3 State of the Art in Sky Models **18**
3.1 Brute Force Simulations 18
3.1.1 The Nishita Model 18
3.1.2 The Bruneton Model 19
3.1.3 The Haber Model 19
3.2 Fitted Sky Dome Models 20
3.2.1 The CIE and Perez Models 20
3.2.2 The Preetham Model 21
3.2.3 Sky Dome Polarization Patterns 22
3.3 Radiance Pattern Captures 22

4 A Full Spectral Analytic Sky-dome Model **24**
4.1 Introduction 24
4.2 Sky Dome Reference Data Generation 24
4.2.1 Atmosphere Composition 24
4.2.2 Describing the Atmosphere via Turbidity 25
4.2.3 Path Tracer Implementation Details 26
4.3 Reference Data: Simulation Results 29
4.3.1 The Effect of Ground Albedo 29
4.3.2 Sunset Color 29
4.4 An Improved Analytical Skylight Model 30
4.4.1 Extending the Perez Model 32
4.4.2 Exclusion of Dusk Conditions 34
4.4.3 Separate Fitting for Reference Data Channels 34
4.4.4 Radiance Distribution Parameters 35
4.4.5 $L_{M\lambda}$ calculation 36
4.5 Error Analysis 36
4.5.1 Comparison with the Original Perez Function 37
4.5.2 Color Renderings, Worst Case Scenario 38
4.6 Direct Solar Radiance 38
4.6.1 Limb Darkening 39
4.6.2 The Analytic Model 40
4.6.3 Model Usage in Practice 42
4.7 Results 43

4.7.1 Skydome Model 43
4.7.2 Solar Radiance Model 43

5 A Full Spectral Analytic Sky-dome model, MK2 48
5.1 Atmospheric Composition Model 49
5.1.1 Synthetic Atmospheric Profiles 49
5.1.2 The Ozone Layer 51
5.2 The Brute Force Simulation 52
5.2.1 The Dataset and Views used for Model Fitting 52
5.2.2 The Atmospheric Path Tracer 53
5.3 The Analytical In-scattered Radiance Model 58
5.3.1 The Dead Ends 59
5.3.2 The Final Model 63
5.3.3 Using the Analytic Model 70
5.3.4 Accompanying Models 70
5.4 Discussion and Results 72
5.4.1 Comparison to Existing Models 72
5.4.2 Aerial Perspective 73
5.4.3 Polarization 73

6 Conclusion 88

Acknowledgements 89

Bibliography 90

List of publications 94

A Attachments 95
A.1 The Atmosphere Description Format (ADF) 95
A.1.1 File Header 95
A.1.2 ADF Blocks 95
A.2 Atmospheric Transmittance 98
A.2.1 Finite In-scattered Radiance Calculation 100
A.3 Error Analysis 101

1. Introduction

The goal of this thesis is to provide a comprehensive overview of the theoretical background and specific techniques that can be used to model the sky and atmosphere in computer generated 3D images, and to present several improvements which were made to the state of the art in this area.

The visual effects caused by the air mass which is present in outdoor scenes are many, some subtle and some not so subtle: the most obvious one is that whenever you look up in the outdoor scene, you see the sky. The sky together with the sun also usually serve as the most important light source in the scene. A subtler, yet very important effect is that the air mass that fills the scene makes distant objects look hazy. We will be henceforth referring to any set of rendering algorithms that tries to replicate some or all of the aforementioned effects as an *atmospheric model* and a model or its part that deals with the skydome specifically as *skydome model*.

Such models have many applications. A good atmosphere model increases the perceived photo-realism of rendered images. We humans are very familiar with the appearance of the sky – it is one of the most commonly seen things in our lives and also commonly visible in rendered scenes. The appearance of the sky gives us certain cues – e.g. a clear sky in the morning has a certain glow around the sun. This glow is caused by the accumulation of water droplets which gradually disappear as the sun goes up and heats the air. Without this effect, the scene does not really have a morning feel and as such, replicating it increases the perceived realism and provides a powerful artistic tool to set the mood of the scene.

The sky also serves as a light source and the illumination varies significantly in intensity and color between noon and twilight, clear and overcast sky etc. The fact that distant objects appear hazy (an effect known as *atmospheric perspective*) allows us to guess their distance and the effect also varies significantly with sky configuration.

In this work we differentiate between photo-realistic, predictive and physically-based models and offer the following definitions of these terms, tailored specifically for the purpose of describing atmospheric models:

A *physically-based* model calculates the results by implementing the equations derived from the laws of physics. The outputs of such a model are physical quantities given in physical units (e.g. radiance in $[W \cdot sr^{-1} \cdot m^{-2}]$). Such a model is *a priori* expected to produce accurate results within the limitations of our current understanding of the laws of light transport and the quality of input data (which in this case would be the vertical profiles of atmosphere composition, scatterer distribution, etc.)

A *predictive* model should also produce radiometrically correct results but it does not necessarily need to arrive to them using physically-based formulas. A model obtained through nonlinear optimization or some other form of example-based machine learning can be considered predictive and can produce accurate results even if it is internally based on an ad-hoc formula, provided it can be guaranteed to yield radiometrically correct results for the entire input parameter space.

For a model to be considered *photo-realistic*, it merely needs to look convincing

to the users. Photo-realism is therefore a very subjective category. Such a model does not need to produce radiometric outputs - an 8bit sRGB image is fine as long as the users consider it to look sufficiently enough like a real sky.

Ideally, we would want an atmospheric model that is physically-based. The advantage of such a model is that its correctness can, if need be, traced back to first principles of currently known radiative transfer principles. In current state of the art, evaluating such a model is however so slow as to be prohibitively expensive even for most offline applications. Therefore we also explore ways to build a model that trades away the physical foundation (but hopefully not the predictive quality) for evaluation speed.

1.1 Our approach

Our work has been heavily influenced by the work of Preetham et al. [1999], the so called Preetham model. We recognize that because the model is easy to implement and integrate, very light on computational resources and last but not least, because it was published with the full set of data required to make the model work makes it a very practical and popular solution.

We started with a goal to produce a model that is similarly easy to use and improves the known shortcomings of the Preetham model. The first step was to obtain high quality reference data. To this effect, we developed a specialized physically-based path-tracer that produces reference spectral renderings of the skydome. We then developed an improved radiance distribution function and an improved meta-model that provides the parameters to the distribution function. We have also developed a matching direct solar radiance function. The results were published together with a reference ANSI C implementation and the complete dataset required to run the model, in the following publications:

- Hošek, Lukáš and Wilkie, Alexander, *An Analytic Model for Full Spectral Sky-dome Radiance.* ACM Transactions on Graphics (SIGGRAPH 2012), 31(4), July 2012.
- Hošek, Lukáš and Wilkie, Alexander, *Adding a Solar-Radiance Function to the Hošek-Wilkie Skylight Model.* IEEE Computer Graphics and Applications, 33(4), May 2013.

Chapter 4 is a significantly extended version of the research published in these publications.

Chapter 5 describes a further development of the skydome model. Significant improvements have been made to the reference path-tracer: the improved version is polarization-enabled and provides a level of physical accuracy that has to date been seen only in specialized optics simulation packages such as `libradtran` [Emde et al., 2016]. The improved reference path-tracer also works with arbitrary atmosphere specifications and can be fed with complex atmospheric composition data such as the aerosol profiles provided by OPAC [Hess et al., 1998]. A new analytical model based on tensor decomposition is presented. This analytical model can fit the reference data to an arbitrary precision and also produces polarization data. It also provides in-scattering data for below-horizon

viewing directions and the camera altitude is now configurable by the user. The new skydome model also comes with an accompanying analytical transmittance function. The model is as of yet unpublished, but finished and tested. It will be submitted for publication once the missing reference data has been computed.

2. Light transport

2.1 The rendering equation

The rendering equation, as stated by Kajiya [1986], reformulated in terms of radiometric quantities is

$$L(x,v) = \epsilon(x_0,-v) + \int_\Omega L(x_0,v') \cdot BRDF(x_0,v',-v)dv' \tag{2.1}$$

where:

- $L(x,v)$ is the radiance $[\mathrm{W}\cdot\mathrm{sr}^{-1}\cdot\mathrm{m}^{-2}]$ reaching point x from direction v,
- x_0 is the point where the ray $x + t\cdot v$ intersects the scene,
- $\epsilon(x,-v)$ is the directly emitted radiance at point x in the direction of $-v$
- $BRDF(x_0,v',-v)$ is the *bidirectional reflectance distribution function* $[\mathrm{sr}^{-1}]$, denoting the ratio of light hitting the surface at x_0 from the direction v' reflected in the direction of $-v$,
- Ω denotes the set of all directions on a hemisphere surrounding x_0

This formulation treats the space in between the surfaces as vacuum, i.e. there is no interaction with participating media going on, and therefore the equation will have to be modified for the purpose of evaluating atmospheric models. But before we do that it is worth mentioning a few words about the evaluation strategies of this formula.

Since the formula is impossible to evaluate analytically for anything but the simplest scenes, it has to be either simplified, approximated numerically, or both.

A Monte Carlo evaluation strategy, where each integral is approximated using one sample, effectively generating a light path through the scene, hence called *path tracing*, has been introduced by Kajiya [1986] along with the rendering equation. Later work on path tracing includes the introduction of *bidirectional path tracing* [Lafortune and Willems, 1996], *multiple importance sampling* [Veach and Guibas, 1995] and *Metropolis light transport* [Veach and Guibas, 1997]. The Metropolis light transport algorithm has also been extended to work with participating media [Pauly et al., 2000].

Replacing the integral in the equation (2.1) with a linear operator M, we get:

$$\begin{aligned} L &= \epsilon + ML \\ (1-M)L &= \epsilon \\ L &= (1-M)^{-1}\epsilon \\ &= \epsilon + M\epsilon + MM\epsilon + M^3\epsilon + ... \end{aligned} \tag{2.2}$$

The intuitive understanding of this expansion is that the total radiance is a sum of a direct term, a once scattered term, a twice scattered term, etc... A simplification, named in the Kajiya paper as the *Utah approximation* omits all bounces except for the first one:

$$L = \epsilon + M\epsilon \tag{2.3}$$

An analogue atmospheric model that only includes light scattered at most once would be called a *single scattering model*. A solution that includes all bounces is a *global illumination solution* and its atmospheric equivalent would be called a *multiple scattering model*.

2.2 Participating media

We consider three separate classes of atmospheric constituents: scattering air molecules, absorbents and aerosols. One particular class of constituents that is not considered in this work are ice crystals. Including these would allow us to simulate atmospheric halos, however that is a quite rare effect not usually required of atmospheric models. A good resource on the topic of the optics of atmospheric halos and their simulation is the HaloSim software website [Cowley and Schoeder, 2004].

2.2.1 Rayleigh scattering

Rayleigh scattering is an interaction of electromagnetic waves with air molecules, named after Johm Willam Strutt, 3rd Baron Rayleigh, who first described it and showed [Strutt, 1899] that the interaction follows from Maxwell's equations.

As light travels through the air mass, the oscillating electric field of the light wave acts on the particles and causes them to become small radiating dipoles. The radiation from these dipoles is the scattered light. The gaseous constituents of air (O_2, O_3, O_4, NO_2, CO_2 etc.) have an effective diameter significantly smaller than the wavelength of the visible light. For particles that are so small relative to the radiation wavelength, the light re-radiates at approximately the same phase over the whole surface. The interaction is elastic, meaning that all energy is re-radiated as light of the original wavelength, scattered in a different direction. The Rayleigh's solution to Maxwell's equation is essentially an analytic solution operating under the simplifying assumption that the whole particle surface re-radiates light at exactly the same phase.

It is worth noting that Rayleigh scattering does not explain the interaction of light with molecular gasses in its entirety – it only explains the elastic interaction component. Notably, a photolytic interaction splits O_2 and O_3 into excited oxygen atoms and a portion of near-UV light is absorbed in this process.

The *cross-sectional area* [m^2] of Rayleigh scattering can be described by

$$\sigma_s^{Ray} = \frac{128\pi^5\alpha_0^2(6+3\rho_n)}{3\lambda^4(6-7\rho_n)}, \tag{2.4}$$

where:

- α_0 is the *polarizability volume* [m^3] (about $1.7 \cdot 10^{-24}$ cm^3 for the mixture of molecules in terrestrial atmosphere [Mclinden, 1998])
- ρ_n is the *depolarization factor* [dimensionless] (0.0279 for naturally polarized light [Mclinden, 1998])
- λ is the *wavelength* [m]

The cross-sectional area (in physics literature also referred to as *macroscopic cross-section*) relates to the probability of a ray of light undergoing a scattering interaction with a single particle and can be thought of as a size of a virtual target that the ray of light has to hit to undergo scattering.

Note the dependence on the 4th power of the wavelength: light at the blue end of the visible spectrum (400 nm) is about 10.5× more likely to be scattered than the light at the red end of the spectrum (720 nm). This is the primary reason why the sky is blue – blue light has a much higher probability of being in-scattered towards the observer's eye. It is also the reason why the solar disc appears to be red at sunset – of the direct light coming from the sun towards the observer, blue light has a higher probability of being out-scattered, and this effect becomes more pronounced at sunset when there is a thicker air mass between the sun and the observer.

When a scattering event occurs, the light is re-radiated in a directional pattern given by the *phase function* [sr^{-1}]:

$$f_{Ray}(\theta) = \frac{3}{4}(1 + \cos^2(\theta)), \tag{2.5}$$

where θ [rad] is the angle between the direction of the incoming and re-radiated light.

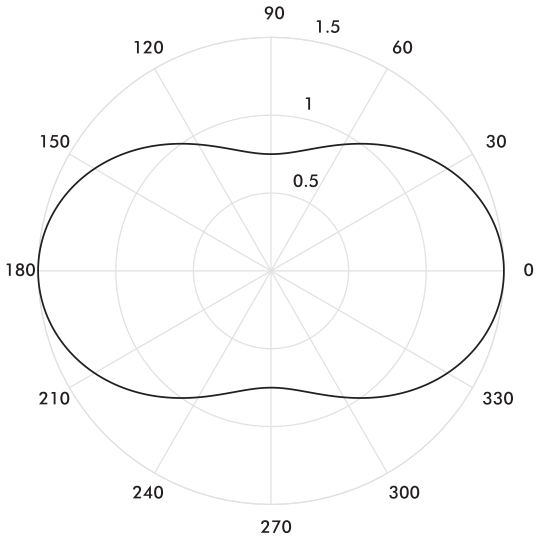

Figure 2.1: The phase function of Rayleigh scattering.

Rayleigh scattering is the dominant scattering in the visible part of the spectrum. A hypothetical atmosphere consisting only of Rayleigh scatterers would already look very similar to a real-world sky on a clear day. Given the shape of the phase function, we can predict that the sky is going to be brightest directly around the sun (at the *solar point*) and at a point directly opposite of it (the *anti-solar point*) and weakest at 90° away from the sun, which is generally true in real world observations.

Polarization characteristics

Light radiated from the sun is randomly polarized, but as it undergoes scattering interactions, each one changes its polarization state. This can be expressed using

Mueller calculus [Mueller, 1948], in which the full state of light is described by a four-component vector (the *Stokes vector*) and the polarizing property of an optical element is described by a *Mueller matrix*.

In a non-polarized path tracer, the state of the light traveling along one ray would be described simply by its radiance. For polarized light, we also need a *reference frame*, given by two vectors $\overrightarrow{x}$, $\overrightarrow{y}$ that are perpendicular to the light direction $\overrightarrow{z}$ and to each other. Given this reference frame, the Stokes vector S describes the polarization state of the light:

$$\overrightarrow{S} = \begin{pmatrix} I \\ Q \\ U \\ V \end{pmatrix} \tag{2.6}$$

where:

- I is the total intensity, and would be equivalent to the radiance in a polarization-unaware path tracer. A Stokes vector representing unpolarized light would have all components except I equal to zero.
- Q is the amount of horizontal-vertical polarization (i.e. amount of polarization aligned with either $\overrightarrow{x}$ or $\overrightarrow{y}$).
- U is the amount of diagonal polarization.
- V is the amount of circular polarization.

A Mueller matrix is a 4×4 matrix and is associated with an *input reference frame* given by two vectors $\overrightarrow{x}$, $\overrightarrow{y}$ perpendicular to each other and to the light direction $\overrightarrow{z}$ and an *output reference frame* $\overrightarrow{x'}$, $\overrightarrow{y'}$, $\overrightarrow{z'}$.

The Mueller matrix of Rayleigh scattering is:

$$F_{Ray}(\theta) = \frac{3}{4} \cdot \frac{1-\rho_n}{1+\rho_n/2} \begin{bmatrix} \frac{1+\rho_n}{1-\rho_n} + cos^2(\theta) & -sin^2(\theta) & 0 & 0 \\ -sin^2(\theta) & 1 + cos^2(\theta) & 0 & 0 \\ 0 & 0 & 2cos(\theta) & 0 \\ 0 & 0 & 0 & \frac{2-4\rho_n}{1-\rho_n} cos(\theta) \end{bmatrix} \tag{2.7}$$

The reference frames for this matrix are:

- $\overrightarrow{z}$ and $\overrightarrow{z'}$ are the incoming and outgoing directions.
- $\overrightarrow{x} = \overrightarrow{x'} = \overrightarrow{z} \times \overrightarrow{z'}$ are the same - a vector picked to be perpendicular to both $\overrightarrow{z}$ and $\overrightarrow{z'}$. This is the normal of the scattering plane.
- The remaining vectors are orthogonal to the rest of the reference frame: $\overrightarrow{y} = \overrightarrow{x} \times \overrightarrow{z}$ and $\overrightarrow{y'} = \overrightarrow{x'} \times \overrightarrow{z'}$.

Since the depolarization factor ρ_n is close to zero for all atmospheric constituents, a simplifying assumption is usually made that it is exactly zero (i.e. the the molecules are completely isotropic) for the purpose of calculating the Mueller matrix. That gives us a formula where the top left component of the

matrix is equal to the scalar version of the phase function presented in equation 2.5.

Looking at the Mueller matrix for Rayleigh scattering, we can predict that the polarization will be tangential to a circle centered around the sun and strongest at 90° away from it. This is indeed true, the real world observations confirm these predictions [Meyer-Rochow and Wehner, 2001, Zhao et al., 2018]. Even though multiple scattering tends to reduce polarization, we can see a degree of polarization of up to 70% on very clear days. Real world observations also reveal two additional features of the sky. One: the sky directly opposite of the sun is not completely unpolarized, but weakly vertically polarized. Two: there is a point lying in between the solar and anti-solar point where the polarization switches from horizontal to vertical, called the Arago neutral point. Our simulation should ideally be able to predict its position.

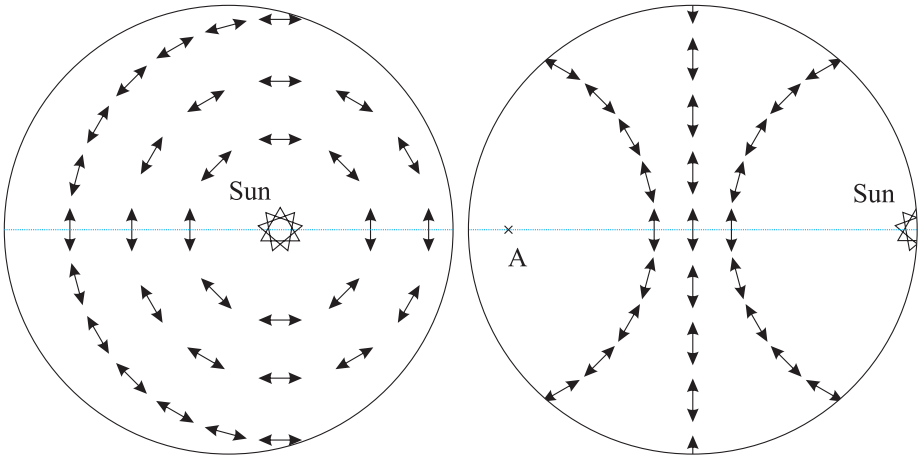

Figure 2.2: Polarization patterns of the sky. **Left:** when the sun is close to the zenith, polarization is mostly horizontal. **Right:** when the sun is setting, polarization is mostly vertical. A marks the approximate location of the Arago neutral point.

2.2.2 Lorenz-Mie scattering

Lorenz-Mie theory [Maheu et al., 1988] gives us a tool to predict the contribution of particles of sizes comparable to the wavelength of visible light. These include water droplets, soot particles, pollen and other atmospheric pollutants generally in the $1\mu m$ - $1mm$ range. For particles of this size we can not assume that the light radiated by the particle acting as a dipole is completely in phase and a geometric optics solution is insufficient because it omits the phase interference effects.

The phase function of Lorenz-Mie scattering is strongly anisotropic and strongly favors light scattered in the forward direction.

In the following section, we will be using these symbols:

- $k = 2\pi/\lambda$ is the angular wavenumber of the light.
- r is the particle radius and $x = kr$ is its size parameter.

- m is the refractive index of the scatterer relative to the index of the ambient medium.

In Jones calculus, which deals with monochromatic plane waves, the polarization property of an optical element is given as a 2×2 *Jones matrix*. For spherical particles, symmetry dictates that the matrix be diagonal. Mie theory gives the following equations for the two nonzero elements of the matrix:

$$\begin{aligned} S_1(\theta) &= \sum_{n=1}^{\infty} \frac{2n+1}{n(n+1)} \left(a_n \pi_n(\cos\theta) + b_n \pi_n(\cos\theta)\right) \\ S_2(\theta) &= \sum_{n=1}^{\infty} \frac{2n+1}{n(n+1)} \left(a_n \tau_n(\cos\theta) + b_n \tau_n(\cos\theta)\right) \end{aligned} \tag{2.8}$$

The functions π_n and τ_n describe the angular scattering patterns of the spherical harmonics used to describe S_1 and S_2:

$$\pi_n(cos\theta) = \frac{P_n^1(\cos\theta)}{\sin\theta} \qquad \tau_n(\cos\theta) = \frac{d}{d\theta} P_n^1(\cos\theta) \tag{2.9}$$

The functions a_n and b_n are the coefficients that control the amplitude of the scattererd field (we use a simplified version that assumes the magnetic permeability of the scatterers and the surrounding medium is the same since air, water, carbon and all the other particle materials that we consider have almost exactly the same permeability as vacuum; the full equations would have to be used for e.g. iron particles). For the full derivation see An Introduction to Atmospheric Radiation [Yung, 2003].

$$\begin{aligned} a_n &= \frac{m^2 j_n(mx) \cdot (x j_n(x))' - j_n(x)(m x j_n(mx))'}{m^2 j_n(mx)(h_n^{(1)}(x))' - h_n^{(1)}(x)(m x j_n(mx))'} \\ b_n &= \frac{j_n(mx)(x j_n(x))' - j_n(x)(m x j_n(mx))'}{j_n(mx)(x h_n^{(1)}(x))' - h_n^{(1)}(x)(m x j_n(mx))'} \end{aligned} \tag{2.10}$$

- $j_n(x)$ and $h_n^{(1)}(x)$ are the spherical Bessel functions.
- P_n^1 are the Legendre polynomials.

The Mueller matrix is then:

$$\begin{pmatrix} I' \\ Q' \\ U' \\ V' \end{pmatrix} = \begin{bmatrix} M_{11} & M_{12} & 0 & 0 \\ M_{12} & M_{11} & 0 & 0 \\ 0 & 0 & M_{33} & -M_{34} \\ 0 & 0 & M_{34} & M_{33} \end{bmatrix} \begin{pmatrix} I \\ Q \\ U \\ V \end{pmatrix} \tag{2.11}$$

where:

$$M_{11} = \frac{S_1 S_1^* + S_2 S_2^*}{2k^2 r^2} \qquad M_{12} = \frac{S_2 S_2^* - S_1 S_1^*}{2k^2 r^2} \tag{2.12}$$

$$M_{33} = \frac{S_2 S_1^* + S_1 S_2^*}{2k^2 r^2} \qquad M_{34} = i\frac{S1 S_1^* - S1 S_2^*}{2k^2 r^2} \tag{2.13}$$

The scattering efficiency and cross-sectional area is:

$$Q_S = \frac{2}{x^2}\sum_{n=1}^{\infty}(2n+1)(|a_n|^2+|b_n|^2) \qquad \sigma_S = \pi r^2 \cdot Q_S \tag{2.14}$$

Mie scattering is also elastic (i.e. the light is re-radiated at the same wavelength), but some energy is absorbed due to ohmic losses. The total extinction can be calculated as follows:

$$Q_E = \frac{2}{x^2}\sum_{n=1}^{\infty}(2n+1)\mathrm{Re}(a_n+b_n) \qquad \sigma_E = \pi r^2 \cdot Q_E \tag{2.15}$$

In a non-polarizing renderer the normalized phase function is:

$$f_{Mie}(\theta) = \frac{M_{11}(\theta)}{\int_0^{\pi} M_{11}(\theta)\sin\theta d\theta}. \tag{2.16}$$

Unlike Rayleigh scattering, where the phase function always has the same shape, the pattern of Lorenz-Mie scattering and its cross-sectional area varies with particle size and wavelength. A further difference to Rayleigh scattering is that the cross-section of Lorenz-Mie scattering is not strongly wavelength dependent and therefore does not give the sky any color tint – a very turbid sky with a high concentration of water droplets is mostly grey.

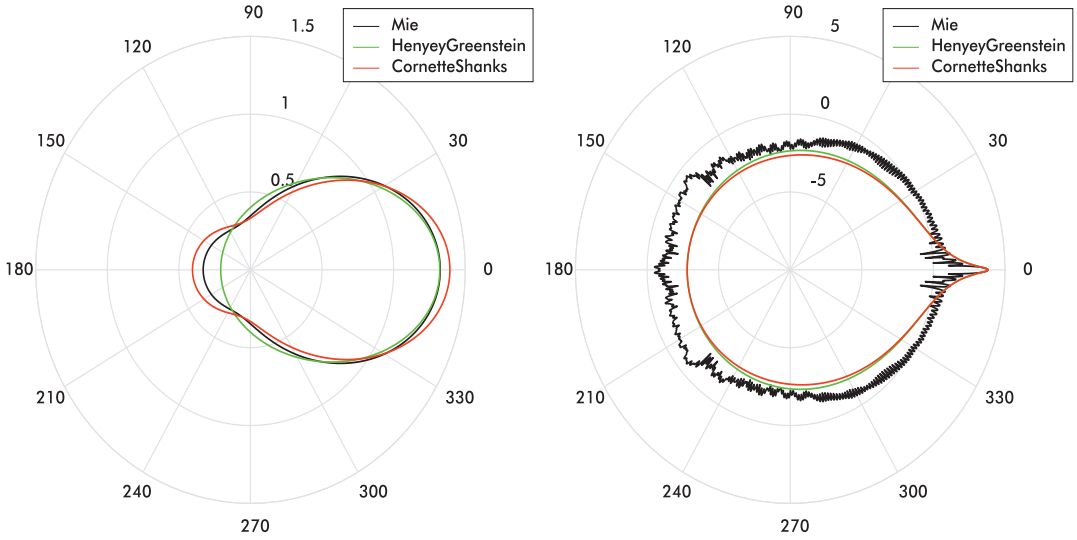

Figure 2.3: The Mie phase function and its approximations. **Left:** a small particle close to the Rayleigh limit, linear plot, $m = 1.33$, $x = 1.2$. **Right:** A large particle, log_{10} plot, $m = 1.33$, $x = 157.07$.

Looking at the plot 2.3, we can see that the phase function has a significant forward bias. The forward facing peak gets narrower and more intense as the size of the particle increases and explains why can can observe a bright glowing area surrounding the sun (called the *circumsolar ring* or *aureole*). Just as one would expect, the intensity of the aureole varies with the concentration of Mie scatterers: on dry and clear days the aureole is subtle, but it becomes very pronounced when the air is moist or dusty. The right figure, which is representative of a particle size found in fog, reveals two additional peaks: a back-scattering peak directly at

180° which produces an optical phenomenon called *glory*, and a peak at around 140° / 220°, which produces a phenomenon called *fogbow*.

Unfortunately, Mie theory does not give us the nice set of closed form formulas that we would like to see in a practical path tracer implementation. Both the cross-sectional area and the phase function can only be expressed in the form of infinite expansion series. We can deal with this problem in a few ways:

The functions can be tabulated. We can tabulate the phase function (and its inverse cumulative distribution function which we need for importance sampling the scattering direction). Tabulated functions are quick to evaluate. The problem is that each wavelength / particle size / particle refraction index combination requires its own table of which we can only have a limited number, but the wavelength and particle size spectrum is continuous and the function changes rapidly with regards to these parameters, so the linear interpolation would not produce meaningful results unless the tabulation would be impractically dense.

The other option is to replace the first-principle formulas by closed-form approximations. One such possible approximation is the formula by Henyey and Greenstein [1941]:

$$f_{HG}(\theta) = \frac{1 - g^2}{(1 + g^2 - 2g\cos\theta)^{3/2}} \tag{2.17}$$

Another approximation was devised by Cornette and Shanks [1992]:

$$f_{CS}(\theta) = \frac{3}{4\pi} \cdot \frac{1 - g^2}{2 + g^2} \cdot \frac{1 + \cos^2(\theta)}{(1 + g^2 - 2g \cdot \cos(\theta))^{\frac{3}{2}}} \tag{2.18}$$

Both feature an *anisotropy parameter* g [dimensionless], which is the average of the cosine of the scattering angle. Provided that we know what the average is for the physically-based phase function, we can plug it into these equations and get a reasonable approximation. Regarding which of these functions to choose, there are a few things to consider:

The Cornette-Shanks function gives a better approximation, particularly for smaller particles. For $g = 0$, the function produces exactly the phase function of Rayleigh scattering (see eq. 2.5), which is consistent with the fact Rayleigh scattering is the limit case of Lorenz-Mie scattering for infinitely small particles.

The downside is that the inverse cumulative distribution function of Cornette-Shanks can not be easily expressed. While there is a closed-form formula, it is impractical to evaluate, and for this reason practically all path tracers use the Henyey-Greenstein version.

A disadvantage of both these formulas is that they provide just the scalar phase function, sans the polarization.

Regarding the approximation quality, fig. 2.3 shows that both functions approximate the forward peak well. Cornette-Shanks is better at reproducing the back-scattering of small particles. Neither of them can however reproduce the secondary peaks that form the glory and the fogbow. The large particle plot also reveals some oscillation around the primary peak. This would suggest that the aureole around the sun should have a colorful, rainbow-like appearance. This would indeed be true if all the scattering particles were of the same size, however since the size is a continuous spectrum and the positions of the oscillation peaks

Figure 2.4: The fogbow and the glory, formed by the secondary peaks of the Mie scattering function. Provided by courtesy of Mila Zinkova.

change with particle size, this colorful effect mostly averages out into neutral white.

2.2.3 Atmospheric composition

To populate an atmospheric model, we need to know the composition of the atmosphere. We can decompose the atmosphere into a finite number of constituent groups: N_2, O_2, water soluble Mie scatterers, water insoluble Mie scatterers, soot, ozone etc. For each of these groups, we need to know:

- The *number concentration* [m^{-3}] of the scatterers as a function of altitude (commonly called *vertical profile*, occasionally also given in [$mol \cdot m^{-3}$])
- In case of Rayleigh scatterers, the *scattering cross-sectional area* [m^2] or alternatively the parameters from which the cross-sectional area can be computed (see eq. 2.4).
- In case of Mie scatterers, the proportional distribution of scattering cross-sectional areas and phase function parameters. This can also be given as an optical index and distribution of particle sizes, from which the other parameters can be computed.
- In case of absorbers, the *absorption cross-sectional area* [m^2].

A very basic atmosphere composition profile can be obtained from the ISO Standard Atomsphere Model [ISO 2533:1975] and commonly available data. According to the ISO model, the atmospheric pressure at sea level is 1013.25 hPa, temperature at sea level is 15° C and the atmosphere is composed of 78.084 % N_2 and 20.946 % O_2 by volume. Using the equation of state, we can calculate the partial pressures and number concentrations of the gasses. The vertical pressure

gradient can be inferred from hydrostatic balance, giving us an exponentially declining vertical profile:

$$N(h) = N_g \cdot e^{-\frac{h}{H_R}} \tag{2.19}$$

where:

- $N(h)$ is the number concentration at a given altitude above sea level h,
- N_g is the number concentration at sea level,
- H_r is the scale height, which is the height where the number concentration is $1/e$ of its sea level concentration (an equal alternative definition is that this is the total height an atmosphere of the same mass would have if its density was uniform). For the ISO atmoshere $H_r = 7994$ m.

The profile described above contains only Rayleigh scatterers and essentially describes a very clear atmosphere devoid of any pollutants. We can already feed this profile into a simulation and it produces results very close to radiometric measurements performed in Antarctica.

To get an atmospheric model for a more realistic scenario we use data obtained by measurements, e.g. from meteorological probes. We use droplet size distribution measurements by Eldridge [1961].

Another source is the excellent OPAC software package [Hess et al., 1998], which contains a large dataset of atmospheric compositions and vertical profiles, both measured and synthetic, representative of certain common weather conditions. The OPAC package defines:

- Aerosol components (insoluble, water soluble, sulfate droplets and various modes of minerals and sea salts). Each component is associated with a specific particle size distribution (given as parameters of a lognormal distribution) and micro-physical and optical properties (most notably for us: scattering and absorption coefficients and the volume phase function).
- Aerosol types (continental clean, continental polluted, urban, maritime...), given as a mixture of the aforementioned aerosol components. A scale height parameter of the vertical profile is also given for each aerosol type – OPAC assumes exponentially decreasing density with regards to altitude.

The OPAC package essentially gives us everything we need to fully populate the atmospheric model with data.

2.3 Extending the rendering equation to account for participating media

Participating media interact with the light in the scene by the mechanisms of scattering and absorption. To account for these, the rendering equation has to be reformulated.

The *transmittance* [dimensionless] between points x and x_0 is the proportion of light radiated from a solid surface at x_0 that reaches the observing point x:

$$T(x, x_0) = exp\left(-\int_x^{x_0} \beta_i^{(e)}(y)dy\right), \tag{2.20}$$

where:

- $\beta^{(e)}$ is the extinction coefficient[m^{-1}]. Extinction is the sum of scattering and absorption.
- The point $x_0(x, v)$ is the point where the ray $x + t \cdot v$ intersects the solid scene

The radiance coming from point x_0 is:

$$I[L](x_0) = \frac{\alpha(x_0)}{\pi} \int_{2\pi} L(x_0, \omega) \cdot \omega \cdot n(x_0)d\omega, \tag{2.21}$$

where:

- α is the surface albedo.
- $n(x_0)$ is the surface normal at point x_0.

Note that this surface reflection model is simpler than the one presented in the original equation. The BRDF is essentially replaced by a constant function, making all surfaces lambertian. This is an acceptable simplification for our purpose – we do not care about the surface properties of any solid geometry in the scene because in our application the solid geometry is never going to be directly displayed. We just need to model that the Earth surface reflects a portion of the incoming light back. The ratio between the reflected and incoming light is the surface albedo.

The radiance per unit distance [$\mathrm{W} \cdot \mathrm{sr}^{-1} \cdot \mathrm{m}^{-2} \cdot \mathrm{m}^{-1}$] in-scattered at point y into the direction of v is

$$J[L](y, v) = \int_{4\pi} \beta_i^{(s)}(y) \cdot f(v, \omega) \cdot L(y, \omega)d\omega \tag{2.22}$$

- $\beta^{(s)}$ is the scattering coefficient [m^{-1}].
- $f(v, \omega)$ is the phase function [sr^{-1}].

Given these, we can rewrite the rendering equation, which calculates the radiance [$\mathrm{W} \cdot \mathrm{sr}^{-1} \cdot \mathrm{m}^{-2}$] reaching point x from direction v, as follows:

$$L(x, v) = (L_0 + R[L] + S[L])(x, v) \tag{2.23}$$

This equation essentially tells us that the incoming light can be broken down into three contributing components.

L_0 is the direct sunlight contribution:

$$L_0(x, v) = T(x, x_0) \cdot L_{Sun} \tag{2.24}$$

R is the light reflected off the solid surface we are looking at. If we are looking into the sky and our view is not obstructed by any solid geometry, x_0 is considered to lie on the intersection of the viewing ray and the atmosphere boundary (in that case, the radiance coming from that "surface" is zero).

$$R[L](x, v) = T(x, x_0) \cdot I[L](x, v) \tag{2.25}$$

S is the light that is in-scattered into our view along the path between the viewer and the solid surface.

$$S[L](x, v) = \int_x^{x_0} T(x, y) J[L](y, v) dy \tag{2.26}$$

For completeness we add that just like the classic rendering equation can be approximated by omitting all bounces except for the first one, we can also make a *single scattering approximation* by eliminating all the recursive occurrences in L. In some applications this might be a useful tradeoff between realism and performance:

$$L = L_0 + R[L_0] + S[L_0] \tag{2.27}$$

3. State of the Art in Sky Models

Within computer graphics, work on sky models tends to fall into three broad categories: first, outright simulation of light transport in the atmosphere. Second, fitted analytic models that attempt to capture the salient aspects of skies within a certain parameter space. And third, efforts that deal with the capture and measurements of skies, either for direct use as scene illumination, or to verify existing models of the other two categories.

3.1 Brute Force Simulations

Examples of this category, such as the work of Nishita et al. [1993] and Haber et al. [2005], tend to yield excellent results in terms of visual quality, but are usually unsuitable for any purpose except reference computations due to their computational cost. The atmospheric research community has also developed general simulation packages such as `libradtran` [Emde et al., 2016], which are significantly more powerful than the one-off solutions introduced within graphics. But as this atmospheric research software is not exclusively intended for the generation of images in the visual range, and as it has considerably higher requirements with regard to general precision and flexibility than normal graphics software, its performance and handling is even more unsuitable for rendering use. For the atmospheric data that is provided along with the package it is, however, a very valuable source of reference solutions, such as those shown in the work of Wang et al. [2016].

3.1.1 The Nishita Model

The model developed by Nishita et al. [1993] is a widely used model representative of the physically based class of sky models. It assumes the atmosphere to be composed of two constituents: gas molecules and aerosol particles.

The air molecules are not differentiated and modeled as Rayleigh scatterers. The vertical profile is an analytic exponentially declining function(see eq. 2.19).

The aerosol particles are modeled as Mie scatterers. The Cornette-Shanks function is used as the phase approximation. The vertical profile is again assumed to be exponentially declining.

The intended use of the model is to render images of the Earth as viewed from outer space. It is essentially a single scattering model: The light from any solid geometry reaching the camera is attenuated by optical density (see eq. 2.20) and in-scattered light is added (eq. 2.26). The in-scattered light does not account for multiple scattering.

Calculation of the transmittance and in-scattered light is essentially done by ray-marching. The part of the ray that travels in the atmosphere ($[P_a, P_b]$) is divided into segments, and for each segment, the transmittance of direct sunlight reaching that segment ($t(A_x, X)$) and the transmittance between that segment and the viewpoint ($t(Pa, X)$) is calculated, along with with the scattering intensity. Essentially, this is the ray-marching solution to the single-scattering approximation (eq. 2.27). The authors argue that the influence of multiple scattering

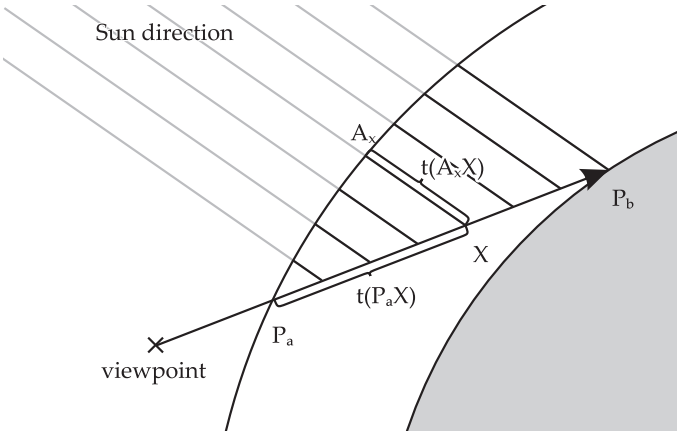

Figure 3.1: In-scattered light calculation in the Nishita model.

and ozone absorption would be negligible.

3.1.2 The Bruneton Model

The model developed by Bruneton and Neyret [2008] works with precomputed data: transmittance between any two points is prepared as a 2D table $\mathbb{T}$. The second is a volume illuminance table $\mathbb{E}$, a 2D table depending on position and sunlight direction. The multiple in-scattering is stored in a 4D table $\mathbb{S}$ depending on position and three angles (view-zenith, sun-zenith and view-sun). When evaluating the model, direct light and single scattering is evaluated by using the precomputed transmittance table. Single scattering can be illuminated piece-wise, occluded and unoccluded ray segments separately, producing the god rays effect. The quality of the model depends of the chosen resolution of the precomputed tables, however for a 12 MB dataset the results are already satisfactory since the contribution of multiple scattering is a very smooth, low frequency information.

3.1.3 The Haber Model

The model developed by Haber et al. [2005] was created with the stated intent to simulate twilight phenomena well. It is similar to the Bruneton model in that it uses precomputed data extensively. The atmosphere is divided into a fixed number of layers. Each of these layers has these fixed optical attributes:

- Scattering coefficient for air molecules (i.e. Rayleigh scatterers),
- Scattering and extinction coefficients and a Henyey-Greenstein anisotropy factor for aerosols (i.e. Mie scatteres),
- Index of refraction.

The two components, Rayleigh and Mie scatterers, are not differentiated any further. The ozone layer absorption is handled by multiplying the unfiltered solar radiation by an absorption factor, which is a function of wavelength and angle of incidence. A unique feature of this model is its ability to handle atmospheric refraction.

The atmosphere layers are divided longitudinally and latitudinally into cells. Each cell stores two values, I_{iso}, an isotropic irradiance value, and I_{aniso}, an anisotropic irradiance value. The precomputation is done iteratively in multiple passes. During the first pass, both I_{iso} and I_{aniso} are initialized from the direct solar irradiance. Each subsequent pass adds one scattering bounce by calculating the energy interchange between each two cells. For the subsequent passes, all irradiance is considered isotropic and is accumulated in I_{iso}.

During evaluation, the ray is traced through the cell array structure, and transmittance and in-scattering is accumulated from every cell the ray passes through.

3.2 Fitted Sky Dome Models

This category has seen quite a number of models being published over the past years [Perez et al., 1993, Preetham et al., 1999]. The models are all based on brute force computations of some sort, and are essentially attempts to fit parametric functions to the properties observed in sky dome radiance simulations of the first category. Common to all these is that they are directly useful for high-quality scene modeling and rendering, and the results of these efforts are nowadays being routinely used in industry. The main drawbacks of these models are typically a more or less limited parameter range, in some cases (particularly for some of the earlier models) a lack of realism for some regions of their parameter space, and a comparatively narrow applicability. For instance, in all existing models, the observer is assumed to be at ground level.

3.2.1 The CIE and Perez 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 standardized 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 color information.

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

$$\mathbb{F}_{Perez}(\theta,\gamma) = (1 + Ae^{B/\cos\theta})(1 + Ce^{D\gamma} + E\cos^2\gamma) \tag{3.1}$$

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

$$\mathbb{F}_{CIE2003}(\theta,\gamma) = (1 + Ae^{B/\cos\theta})(1 + C(e^{D\gamma} - e^{D\frac{\pi}{2}}) + E\cos^2\gamma) \tag{3.2}$$

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 normalized 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* [$\mathrm{cd}\cdot\mathrm{m}^{-2}$]:

$$Y(\theta,\gamma) = \frac{\mathbb{F}(\theta,\gamma)}{\mathbb{F}(0,\theta_S)} \cdot Y_Z \tag{3.3}$$

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 [CIE04], 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.

3.2.2 The Preetham Model

The widely used model developed by 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 optimization, 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, the simple formulas make it suitable to implement it as a shader program, and most importantly the paper comes with all the coefficients required to make the model work.

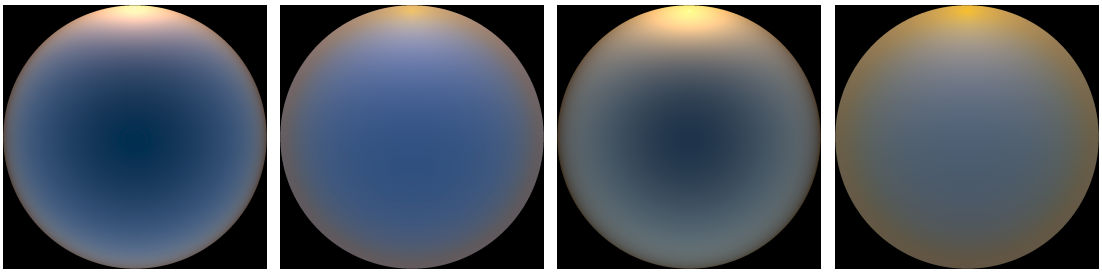

(a) T = 4, reference (b) T = 4, Preetham (c) T = 6, reference (d) T = 6, Preetham

Figure 3.2: Comparison of reference solutions (computed according to the techniques outlined in section 4.2) 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 color pattern and the luminance distribution are considerably off.

The overall "mood" of the sky is controlled by a single turbidity parameter (more on that in section 4.2.2; for now let us just say that the number tells us how clean or hazy the sky is). 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 3.2. 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.

Despite the known limitations, the advantages still make it a 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.

3.2.3 Sky Dome Polarization Patterns

An unusual addition to the line-up of analytical models was the polarization work of Wilkie et al. [2004]. In this paper, an approximative solution to the problem of sky dome polarization patterns was proposed. As the authors did not have access to any reference data, nor to a brute force simulator capable of including the phenomenon, they entirely based their approach on indirect reasoning: and as Wang et al. [2016] conclusively demonstrated, the results were rather sub-optimal. The authors of the latter paper also proposed a new, high-quality analytical model of skydome polarization that can be efficiently evaluated. The only drawback of the proposed new polarization pattern model is that no matching analytical model of skydome radiance is provided. As polarization data is typically only needed for highly accurate renderings that would require matching skydome radiance data, this omission undermines the utility of this promising new model. By contrast, our new proposed model provides an integrated solution that covers all aspects of skydome radiance (in-scattered radiance, attenuation, polarization) from one and the same brute force reference simulation.

3.3 Radiance Pattern Captures

On-set captures of sky-dome radiance for re-lighting purposes have been a standard technique in the movie industry for many years now: collections of representative stock HDR sky images are a common lighting resource in production environments. Both rendering software, as well as production workflows, routinely include such HDR lighting data. However, such production-level HDR captures are rarely, if ever, taken in a calibrated manner that would allow one to use them as either basis for renderings done in absolute units, or as basis for verification of second category models.

A recent and highly useful exception is the work of Kider et al. [2014], who provided a systematically collected dataset of sky dome radiance for representative conditions that can also be used to directly illuminate scenes. The only drawback of that work is that it failed to capture the exact atmospheric parameters at the time of capture: while available meteorological data was meticulously

recorded, no detailed information on the particle distribution profiles in the atmosphere was collected. This lack of scatterer profile data means that it is generally not easily possible to attempt a brute force computation of the first category that matches the captured luminance distributions in the dataset of Kider et al. [2014]. There are too many degrees of freedom in how a clear atmosphere can be structured to be able to derive scatterer profiles from the luminance distributions alone. Which, in turn, makes it difficult to *directly* base any new second category model on these measurements.

4. A Full Spectral Analytic Sky-dome Model

This chapter is based on the following publications:

- *"An Analytic Model for Full Spectral Sky-dome Radiance"*, published in ACM TOG (SIGGRAPH 2012 proceedings) [Hošek and Wilkie, 2012].
- *"Adding a Solar-Radiance Function to the Hošek-Wilkie Skylight Model"*, published in IEEE Computer Graphics & Applications [Hošek and Wilkie, 2013].

The results were published accompanied by ANSI C reference implementations which includes all the data required to run the fitted model.

4.1 Introduction

In this chapter, 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. 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.

4.2 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 an overview of the techniques used by this path tracer and the atmosphere composition model used to obtain the reference data.

4.2.1 Atmosphere Composition

Our reference model considers the atmosphere to consist of two components: air molecules (Rayleigh scatterers) and aerosol particles (Mie scatterers). Both of them are not differentiated any further.

For Rayleigh scatterers, this is a perfectly legal simplification: all Rayleigh scatterers, regardless of their actual chemical composition, have the exact same phase function. Due to this, an atmosphere consisting of multiple Rayleigh scatterers with their respective scattering coefficients $\beta_1, \beta_2, ...\beta_n$ can be effectively replaced by a single Rayleigh scatterer with a scattering coefficient $\beta = \sum_{i=1}^{n} \beta_i$ and the radiative transfer simulation will behave the same.

The vertical profile is a function that provides *number concentration* [m^{-3}] (i.e. the number of particles per unit volume) for a given altitude. For the Rayleigh scatteres, we used the exponential fall-off function (see eq. 2.19).

For Mie scatterers, this gets trickier. We use a Henyey-Greenstein approximation of the phase function, which requires that we supply the anisotropy parameter g. For a single particle composition and size, the parameter would be possible to obtain by the cosine rule. In our simulation, we consider all the particles to have the same composition (water droplets, refractive index $n = 1.33$, i.e. we do not consider soot, pollen or any other water insoluble constituents), but we consider their size to follow a specific distribution (an empirically measured [Eldridge, 1961] fog drop size distribution). To get an anisotropy factor representative of the whole spectrum of particle sizes, we have calculated a weighted average of the Mie phase functions for all sizes, weighted by the scattering coefficients of each size category. The anisotropy factor g was then found by non-linear regression of the Henyey-Greenstein function to the averaged Mie function. Since the reference renderer is a full spectral renderer, this process had to be done separately for each spectral channel.

4.2.2 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, and has already been used in the model by Preetham et al. [1999]. It is defined as the ratio of the additional *optical thickness* [dimensionless] of the atmosphere in question t_h to the optical thickness of an idealized atmosphere that consists only of pure gas t_m:

$$T = \frac{t_m + t_h}{t_m} \tag{4.1}$$

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.

The vertical profile of Mie scatterers is again an exponential fall-off function. Scale height was chosen to be $H_r = 1200$ m. This is a somewhat arbitrary choice based solely on the fact that it is also used by other authors [Elek and Kmoch, 2010]. Number concentration at sea level N_g was computed so that the optical thickness of the atmosphere at sea level conforms to the turbidity equation

(eq. 4.1). By definition, $T = 1$ produces an atmosphere with no Mie scatterers whatsoever.

4.2.3 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 color, and another sphere that represents the upper boundary of the atmosphere, placed 100 km above sea level. In hindsight, the upper boundary is placed excessively high. It should have been placed at 50 km or even lower. When rendering images from the ground level, the contribution from the atmosphere above 50 km is negligible.

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 visualizations, the spectral results were converted to CIE XYZ values and then tone-mapped using the "Interactive Calibration" algorithm proposed by Matkovic and Neumann [1996].

Scattering and Absorption Event Generation

In our case we consider three possible interactions: Rayleigh scattering, Mie scattering and Mie absorption. Each one of these interactions is associated with a corresponding cross-sectional area σ_i $[m^2]$. We also know the number concentrations N_i $[m^{-3}]$. The *scattering / absorption coefficients* $[\mathrm{m}^{-1}]$ are obtained simply as:

$$\beta_i = \sigma_i \cdot N_i \tag{4.2}$$

and represent the mean number of collision events per a ray segment of unit length. The *total extinction coefficient* β of the medium can be calculated as:

$$\beta = \sum_{i=1}^{n} \beta_i \tag{4.3}$$

Assuming that medium is homogeneous (i.e. β is constant), we can derive the formula for generating a *random free-flight path* [m]:

$$p = -\frac{1}{\beta} \cdot logR \tag{4.4}$$

R denotes a random number picked from $[0, 1]$ with an uniform distribution. The term $\frac{1}{\beta}$ is occasionally referred to as *mean free-flight path* [m].

We are now faced with the problem that the atmosphere is not homogeneous. To deal with this we use a technique called *Woodcock tracking* (also referred to as *delta tracking* or *zero collision method*) [Raab et al., 2008]. This method introduces an additional virtual type of collision called the zero collision. Its scattering coefficient is calculated so that the toal scattering coefficient stays constant throughout the whole medium. When this collision is then encountered, instead of the ray being scattered or absorbed, it is simply recast in its original

direction. More formally: the collision coefficients β_i are now functions of position (or, in our case, altitude). Let us redefine β:

$$\beta = \max_h \left(\sum_{i=1}^{n} \beta_i(h) \right) \tag{4.5}$$

β is now the *majorant* of the total collision coefficient. Since we have chosen all the atmospheric constituents to have an exponentially declining density, we know for a fact that the maximum occurs at $h = 0$. At any given altitude, the sum of $\beta_1, ...\beta_n$ can only be less or equal than β. The additional virtual interaction is introduced with its corresponding collision coefficient:

$$\beta_{n+1}(h) = \beta - \sum_{i=1}^{n} \beta_i(h) \tag{4.6}$$

We now have a participating medium with $n+1$ interactions and a constant total extinction coefficient equal to β. When tracing a ray through the atmosphere, a random free flight path is generated first using eq. 4.4. That free flight path is terminated by a collision at an altitude h_{col}. We now have to determine which of the possible $n + 1$ interactions occurred by picking one of them randomly with a probability proportional to their collision coefficient $\beta_i(h_{col})$. If the interaction is Rayleigh or Mie scattering, the ray is recast in a random direction with a probability distribution conforming to the phase function. If the interaction is absorption, the ray is terminated. In case of zero scattering, the ray is recast in its original direction. The method is essentially a form of rejection sampling – if a collision occurs and it is chosen to be a zero collision, that collision is rejected. We can recast the ray in its original direction from the rejected collision point, since the collision probability per traveled path does not depend on events that occurred in the particle's history and any point in the path can be considered the starting point of a new path.

Multiple Suns Optimization

The sun is the only source of light in this setup. The angular diameter of the sun is approximately $32'$ (the value slightly varies since Earth's orbital trajectory is not perfectly circular), which means that the sun only covers about 0.0005% of the total celestial sphere. In a primitive path tracer implementation an overwhelming majority of the rays would be wasted, because they would exit the atmosphere and not hit the sun. General purpose path tracers employ various techniques such as light source sampling, bi-directional path tracing or Metropolis light transport to make the Monte Carlo simulation converge faster. We have chosen a different approach that is both easy to implement and conceptualize and exploits the symmetry of our scene.

Firstly, we can 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. The 10 scenes are completely identical except for the fact that the sun is placed differently. Therefore, the simulation can be run just once. A random ray path is generated as usual and when the ray exits the atmosphere, instead of checking for collision with the sun in just one position, every possible position is being checked. If the ray hits any of the suns, the information on which one it was is

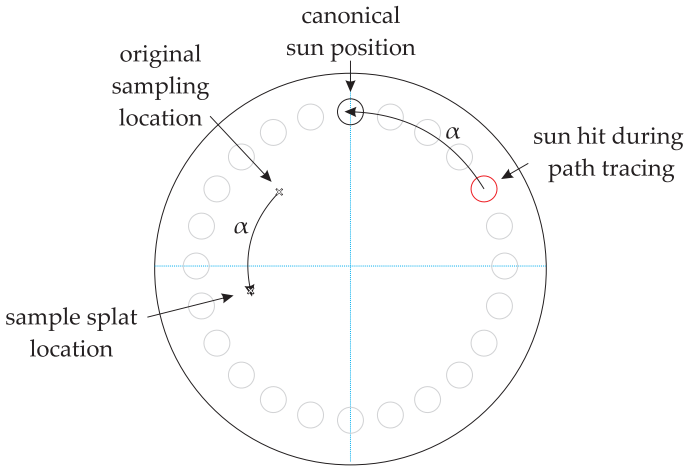

Figure 4.1: The multiple sun approach: In addition to the sun in its original position, there are rotated suns (displayed in grey) at rotated positions, still keeping their original solar elevation and distributed in a circle centered around the zenith. If a ray cast in the original direction hits a rotated sun instead of the canonical one, the final splatting location has to be rotated by the same angle that would be required to rotate the sun to its canonical position.

passed back to the ray sampler and the sample is splatted to the corresponding image. This approach yields a 10× 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 place on the projection plane. When rendering an image with solar elevation 0°, we can e.g. place 720 suns around the horizon without overlapping (since each of them is 0.5° wide) and get up to 720× speedup, 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 already fast enough as a reference solution generator with the described "many suns" optimization. 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.

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 polarizability* 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 analytic 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.

4.3 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 4.2. We systematically performed 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 color of the sunset.

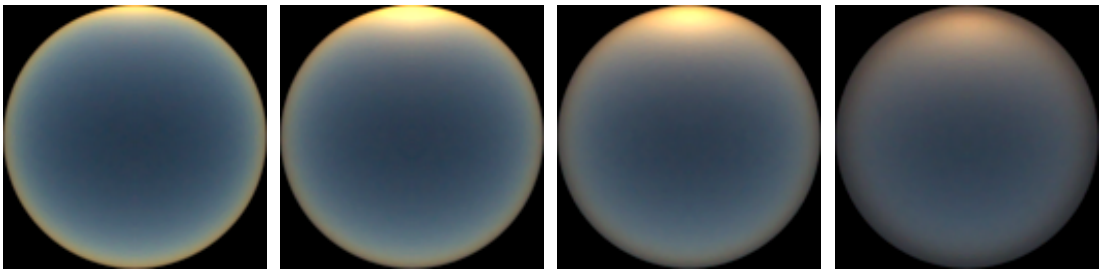

(a) T = 2 (b) T = 4 (c) T = 6 (d) T = 8

Figure 4.2: 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.

4.3.1 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 550 nm, so almost all incident radiation is reflected and back-scattered towards the viewer due to Mie scattering. On the other hand, grass has albedo of about 0.09 at 550 nm. Especially in high turbidity settings, changing ground albedo alters the brightness of the whole sky-dome in a perceptible fashion. Figure 4.3 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 colors, or the better handling of high turbidity situations, emerged as further results once we started work on the modification process. But as figure 4.13 shows, the changes in sky-dome appearance due to ground albedo are not restricted to monochrome brightness changes.

4.3.2 Sunset Color

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 4.4c and 4.4d show a comparison between a fisheye photo of a real sunset sky, and our reference rendering. Figure 3.2 shows side by side comparisons of our analytical model (which is visually indistinguishable

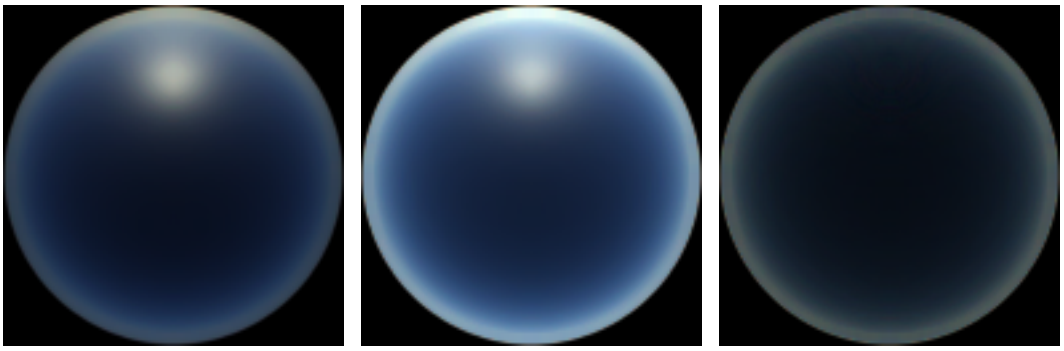

(a) $\alpha = 0.1$ (b) $\alpha = 0.9$ (c) difference

Figure 4.3: Brute force path tracer renderings. Difference between low and high ground albedo α, solar elevation 40°, $T = 4$. Images rendered in linear color space using the same tone mapping setting.

from the reference rendering) to the Preetham model for similar turbidity values, and solar elevations. Both the absolute colors, as well as the distribution of the colors in the Preetham model can only be considered a rough approximation at best.

4.4 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 role (they actually already have a visible impact at turbidity as low as 5, but at the time of the publication we did not consider this to matter as much). We therefore restricted ourselves to a scenario of a gradual increase in haze up to moderate haze levels. A CIEXYZ and RGB 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). Each spectral channel was rendered completely separately and treated as a monochromatic rendering performed exactly at the specified wavelength – the physical quantity produced by the simulation is *spectral radiance* $[\mathrm{W} \cdot \mathrm{sr}^{-1} \cdot \mathrm{m}^{-2} \cdot \mathrm{nm}^{-1}]$. 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 stabilizes and subsequent changes are minimal. To

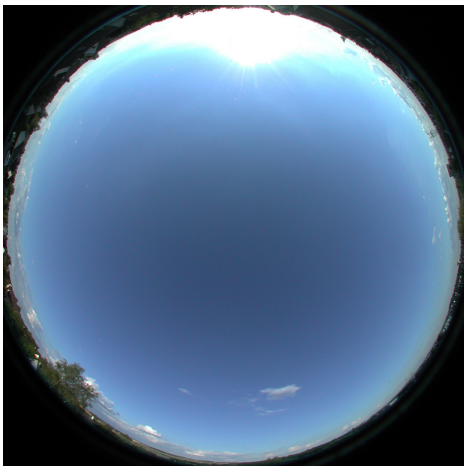

(a) Sunrise

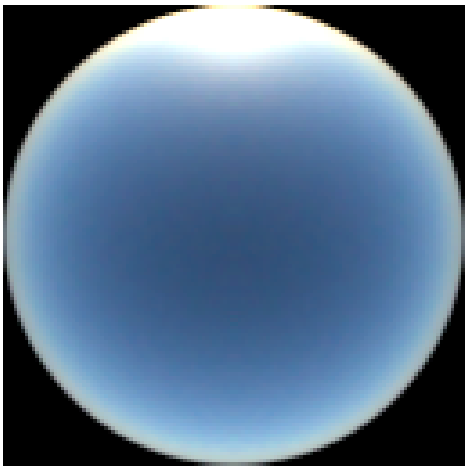

(b) $\alpha = 0.1$ T=3

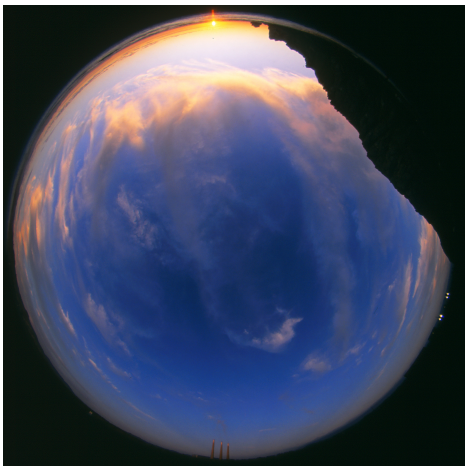

(c) Sunset

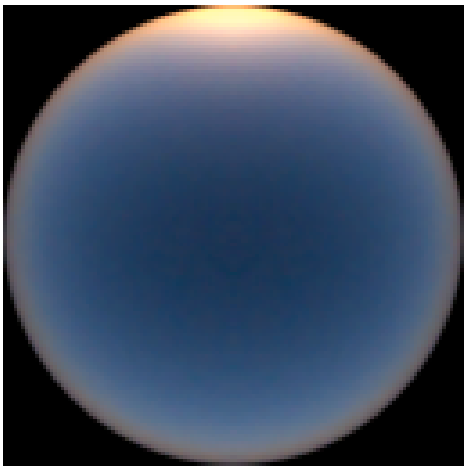

(d) $\alpha = 0.1$ T=4

(e) Daytime sky

(f) $\alpha = 0.9$ T=7

Figure 4.4: Comparison between fisheye photos (left) and brute force path tracer outputs (right). 4.4c was taken on Velvia film, which produces characteristically saturated pictures. To compensate, the path traced image was rendered with a white point of 4400K.

better capture this behavior, the reference images were not generated at equidistant solar elevations. Instead, more samples were taken at lower solar elevations: the n-th image ($n \in \{0, 1 \ldots, 45\}$) was placed at solar elevation $(\frac{n}{45})^3 \cdot 90°$. 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, i.e. there is no directly visible solar disc on the rendered images. This decision was made because including the solar disc 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 skydome model therefore does not include the solar disc. The technique for rendering the disc itself is described in section 4.7.2.

4.4.1 Extending the Perez Model

The Perez formula can capture luminance distributions of skies with low turbidities almost flawlessly, providing the $A, B, \ldots E$ 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 the circumsolar ring phenomenon. This produces a highly localized 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 localized 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 is an obvious reason 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 bias that moves the offending division by zero about 0.5° below horizon and also solves the extremely bright rim around horizon.

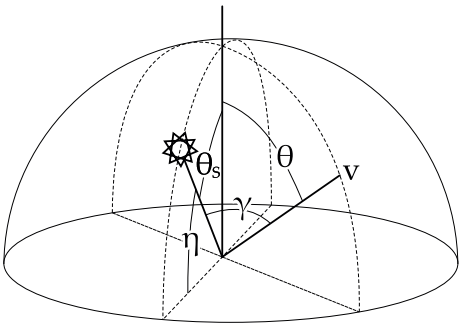

Figure 4.5: The coordinate system used in the model.

The final formula is:

$$\begin{aligned}\mathbb{F}(\theta,\gamma) =&(1+Ae^{\frac{B}{\cos\theta+0.01}})\cdot(C+De^{E\gamma}+ \\ &+F\cos^2\gamma+G\cdot\chi(H,\gamma)+I\cdot\cos^{\frac{1}{2}}\theta)\end{aligned} \tag{4.7}$$

$$\chi(g,\alpha) =\frac{1+\cos^2\alpha}{(1+g^2-2g\cdot\cos\alpha)^{\frac{3}{2}}} \tag{4.8}$$

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 3.1, 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 3.3). Our model can now reproduce the huge spikes in brightness around the sun, which poses a problem with very high solar elevations ($> 85°$), because the function 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

$$L_\lambda = \mathbb{F}(\theta,\gamma)\cdot L_{M\lambda} \tag{4.9}$$

where $L_{M\lambda}$ is the mean luminance value across the upper hemisphere. This implies that the mean 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 4.7.

4.4.2 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 at the time of publication extend the fitting process beneath solar elevations of $0°$. Figure 4.6 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. An extended model that can also replicate after-dusk skydomes, including the Earth's shadow, will be presented in chapter 5.

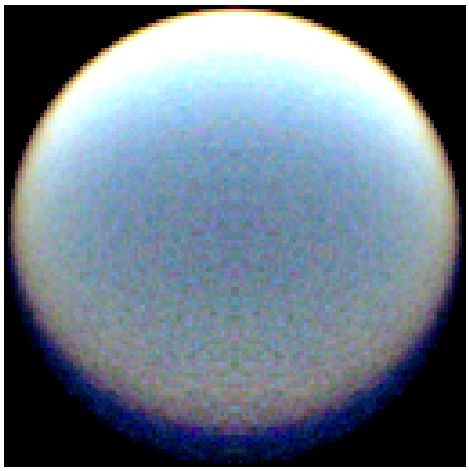

Figure 4.6: An over-exposed reference rendering of the sky at solar elevation $-5°$. The crescent-shaped shadow which the Earth starts to cast into the atmosphere once the sun has gone beneath the horizon can be seen in the lower part of the image.

4.4.3 Separate Fitting for Reference Data Channels

The new model comes in three versions: one fully spectral, one that supplies CIEXYZ values and one that supplies RGB values (in linear space, with sRGB primaries). In all versions, the model produces radiance data that is separately fitted for each channel provided in the reference data, be that spectral values, or the XYZ and RGB versions. The rationale behind the separate fitting for each band is obvious from figure 4.7: no single set of coefficients could properly capture the variations in pattern seen across the spectral range.

The separate fitting for each wavelength band or color channel of course means that when using the model, one has to calculate formula 4.9 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),

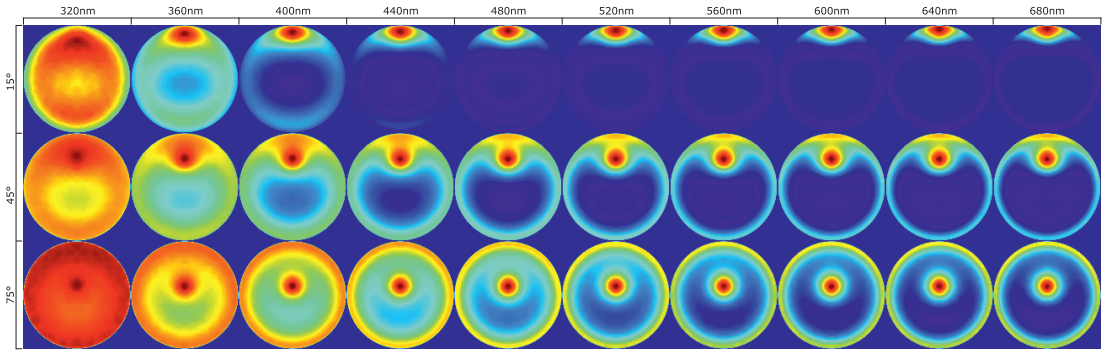

Figure 4.7: Brute force simulation results for the range from UV to the red part of the spectrum for three solar elevations. Separate channels of spectral radiance in false color. Note the substantial changes in the patterns, in particular for the UV range. Fitting all these patterns to a single set 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.

and most importantly, the ability to control ground albedo separately for each spectral band / color channel.

4.4.4 Radiance Distribution Parameters

Radiance distribution parameters are the variables $A, B, \ldots, I$ that plug into equation 4.7. 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 4.7 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 $T \in \{1, \ldots, 10\}$, two albedo values $\alpha \in \{0, 1\}$ and nine distribution parameters $p \in \{1, \ldots, 9\}$ there is a set of six control points $c \in \{1, \ldots, 6\}$. 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:

$$
\begin{aligned}
v_p^\lambda = & m_{T,\alpha,p,1}^\lambda \cdot (1-x).^5 + \\
& m_{T,\alpha,p,2}^\lambda \cdot 5x(1-x).^4 + \\
& m_{T,\alpha,p,3}^\lambda \cdot 10x^2(1-x).^3 + \\
& m_{T,\alpha,p,4}^\lambda \cdot 10x^3(1-x).^2 + \\
& m_{T,\alpha,p,5}^\lambda \cdot 5x^4(1-x).^1 + \\
& m_{T,\alpha,p,6}^\lambda \cdot x^5
\end{aligned}
\tag{4.10}
$$

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 4.7. The arrays were fitted to the outputs of the reference path tracer using Levenberg-Marquardt nonlinear 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 degenerate. To ensure this does not happen, interpolated results were included in the error function during the fitting process.

4.4.5 $L_{M\lambda}$ 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 mean value of radiance that is plugged into formula 4.9. The values in the table were again fitted using Levenberg-Marquardt and intermediate values can be obtained simply by interpolating.

4.5 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 $20\log_{10}$ 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 in Appendix A.3.

	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 4.1: For each waveband/color 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 4.8.

4.5.1 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 4.4.1. It should be noted that the SNR results we show in table 4.1 are for the new model "as is" (i.e. obtained using the above mentioned bezier interpolation), while the parameters for the Perez formula used in this comparison were obtained by fitting the Perez model separately to every single waveband/turbidity/albedo/solar angle in the 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 4.8. 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.

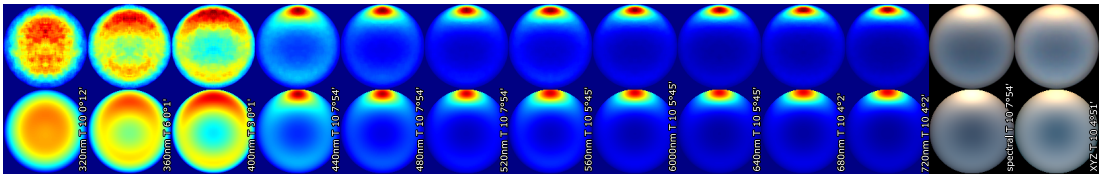

Figure 4.8: The lowest SNR fittings of the new model for each waveband of the spectral model (spectral radiances in false color), as well as for both the overall spectral and CIEXYZ versions of the model (converted and tonemapped to a RGB image). **Top:** reference results. **Bottom:** output of the new model. These images correspond to the lowest SNR cases mentioned in table 4.1. Note that for very short wavelengths, the noise present in the reference data is a major factor in low SNR readings.

4.5.2 Color Renderings, Worst Case Scenario

For this test, spectral model outputs for each turbidity/albedo/solar elevation configuration were rendered, then converted to CIEXYZ via the CIE standard observer color 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 4.8.

4.6 Direct Solar Radiance

As mentioned previously, the path tracer was programmed to discard rays that do not undergo any scattering events along their way – in other words, rays that directly hit the sun are not included in the reference images and as a result, they are not included in the analytic model either. There is a good reason for this: the solar disk creates a very small discontinuous area on the sky-dome that can be up to 18 stops brighter than its surroundings. The discontinuity and brightness difference would pose a problem for the stability of the optimisation process and the radiance pattern of the solar disk exhibits properties that would be hard to capture using the framework used for calculating the sky-dome distribution parameters.

Recall the modified rendering equation presented in eq. 2.23. The appearance of an object as perceived by the observer can be broken down into emitted light (L_0), reflected light (R) and in-scattered light (S). The skydome model provides us with the S component, and for most of the sky, that is all we need, because there are no direct rays coming from the darkness of the empty space. But that changes when the sun is in the way and the purpose of the solar radiance model is to fill the gap represented by the missing $L_0 = T \cdot L_{Sun}$ term.

It is worthwhile to observe that the transmittance T is not dependent on the position of the sun on the sky dome. Given a fixed configuration of atmospheric conditions, transmittance for a fixed ray elevation will stay the same, regardless of solar elevation or ground albedo. The same path tracing routine that was

used for obtaining the reference sky-dome images can be used with only minor modifications to get the estimate of atmospheric transmittance – we can just cast rays in one direction and calculate the ratio of rays that were not scattered before they escaped the atmosphere.

4.6.1 Limb Darkening

L_{Sun} can be obtained from the NASA tabulated data [Coulson, 1975], but this only gives us the spectral radiance at the center of the solar disc. This information is not yet sufficient to generate reasonably accurate renderings of the sun, since the sun also exhibits a property called *limb darkening*. Light emitted from the surface of the sun has to travel through sun's atmosphere, which attenuates light. Rays coming from the edge of the solar disc have to travel a longer path almost parallel to the sun's surface in its atmosphere, which produces an effect that leads to the edges of the solar disc being darker than the center. In astronomy, it is common to use a simple approximate model to describe the phenomenon:

$$\frac{I(\psi)}{I(0)} = \sum_{k=0}^{N} a_k \cos^k(\psi) \tag{4.11}$$

where $I(\psi)$ is the *radiant intensity* $[\mathrm{W}\cdot\mathrm{sr}^{-1}]$ at a point on the solar disc along the ray path forming angle ψ with respect to the stellar radius. $I(0)$ is the radiant intensity at the center point.

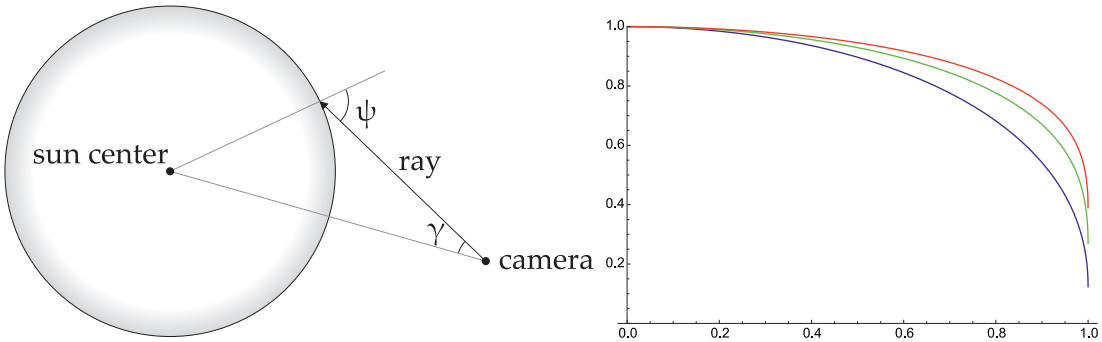

Figure 4.9: **Left:** limb darkening geometry. **Right:** plot of the limb darkening attenuation (vertical axis) for short, medium and long wavelengths across the solar disc radius (horizontal), plotted in blue, red and green, respectively. On the horizontal axis, 0 is the center, and 1 the outer rim of the solar disc. Note the difference in residual brightness at the rim: blue light is attenuated twice as much as red light. This leads to the orange tinge on the rim of the solar disc that can be seen in the top left image of figure 4.10

The coefficients a_k are available for wavelengths in the whole visible as well as near ultraviolet and infrared parts of the spectrum in tabulated form [Neckel and Labs, 1994]. In our model, we use a slightly adapted form of the coefficients found in that paper: we interpolated the coefficients for the exact wavelengths used in the skydome model.

The term $\cos(\psi)$ can be calculated as

$$\cos(\psi) = \sqrt{1 - \frac{\sin^2\gamma}{\sin^2 0.25^\circ}} \tag{4.12}$$

where γ is the same sun center-ray angle used in skydome models (see figure 4.5) and $0.25°$ is sun's angular radius as viewed from Earth.

As figure 4.10 shows, inclusion of limb darkening has a perceptible and pleasantly realistic, if not overly dramatic, visual impact on the appearance of the solar disc. Which might lead to the conclusion that it can be safely omitted for all those images where no close-ups of the solar disc are shown. Which, in practice, means most outdoor scenes illuminated by skylight, since the sun usually only acts as a light source, and is not directly visible. But it is crucial to realize that limb darkening has an additional effect beyond mere realism of sunset images: namely, that it reduces the radiant power of the unattenuated solar disc by approximately 23%, compared to a solar disc of an uniform color equal to the center brightness. This has direct implications for usage of the model presented here, which are discussed at the end of section 4.6.3.

4.6.2 The Analytic Model

Just like the skydome model, the solar radiance model provides independent data for each spectral band – two ultra-violet bands (320 nm and 360 nm) and nine bands in the visible spectrum (400 nm - 720 nm in 40 nm increments).

For each spectral band λ and each integer turbidity value T in the range $[1, 10]$, we provide coefficients for a piece-wise cubic Hermite interpolation polynomial function $r_{\lambda,T}(\epsilon)$ that approximates the $L_0 = L_{Sun} \cdot T$ term of equation 2.23, that is, transmittance for a given ray elevation ϵ multiplied by the solar radiance taken from the NASA dataset (i.e. solar radiance unattenuated by Earth atmosphere or limb darkening). In other words, the function $r(\epsilon)$ gives us the direct solar radiance for a ray shot at elevation ϵ, assuming it would hit the sun center.

The piece-wise cubic Hermite interpolation function $r_{\lambda,T}(\epsilon)$ consists of 45 segments with boundary points x_i placed at

$$x_i = \frac{\pi}{2} \cdot \sqrt[3]{\frac{i}{45}} \tag{4.13}$$

The non-equidistant placement of segment borders was chosen to get a finer resolution for ray elevations close to the horizon, where the solar radiance changes more dramatically and an actual radiance gradient can be seen on the solar disc.

Provided that the ray elevation angle ϵ falls into i-th segment (i.e. between x_{i-1} and x_i), corresponding coefficients $m_{\lambda,T,i,k}$ are read from a table and the value of function $r_{\lambda,T}(\epsilon)$ is calculated as

$$r_{\lambda,T}(\epsilon) = \sum_{k=0}^{3} m_{\lambda,T,i,k} \cdot (\epsilon - x_{i-1})^k \tag{4.14}$$

To get a limb-darkened spectral radiance value, multiply the output of function r by the formula from equation 4.11. If values for non-integer turbidities or wavelengths between the provided channels are required, linear interpolation provides sufficient results.

(a) $T = 1, \eta = 42°$

(b) $T = 1, \eta = 0.5°$

(c) $T = 6, \eta = 0.5°$

(d) $T = 7, \eta = 0.25°$

Figure 4.10: Limb darkening at different turbidity and solar elevation settings. The white solar disc in the top-leftmost image is not an anomaly: at higher elevations, the sun is practically not attenuated, and has an approximate color of D65, which is the white point of most imagery.

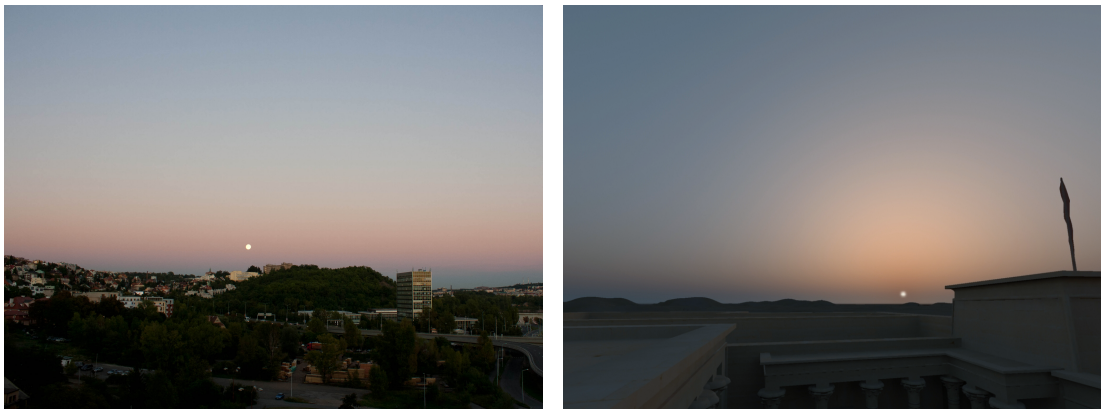

(a) Reference photograph (b) Our model, T=9

Figure 4.11: A hazy sunrise. Note that the new model includes the dark rim that appears around the horizon at higher turbidities.

4.6.3 Model Usage in Practice

It is important to note that as the solar radiance function introduced in the previous section changes with elevation, one has to invoke it separately for the varying elevations found on a given solar disc in order for a gradient to appear.

When using this function, it is also vitally important to not use it on its own, but to add the output of the skydome model to it. The solar radiance function we discuss here only captures the direct illumination, to which the energy from all photons that have scattered at least once has to be added - effectively, the in-scattered light in the area of the solar disc, which is exactly the information that the skydome model provides. A failure to take this into account (i.e. only using the direct solar radiance) results in a much too dark solar disc at lower elevations, especially for higher turbidity situations.

As mentioned in section 4.6.1, due to the non-uniform brightness distribution on the sun that is present even at high solar elevations, care has to be taken when determining the overall radiant power of the solar disc. In practice, a sufficiently large number of random samples on the disc has to be averaged to get a reasonable result. Due to limb darkening, and at lower solar elevations also to the non-linear progression of attenuation, using just a single sample at the centre of the solar disc yields significantly wrong results: on average across wavebands, the error induced by limb darkening alone is about 23%.

Due to the non-linearity of the solar brightness function across elevations and wavelengths, it is not feasible to find a simple linear correction factor for this effect. The only reliable way to determine the energy of the solar disc is to take multiple samples, and to average those. Since this only has to be done once for a given solar elevation, this does not pose a significant run-time overhead, though.

4.7 Results

4.7.1 Skydome Model

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 localized aureoles and sunrise/sunset colors. 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. Furthermore, 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 4.12 demonstrates the impact that inclusion of UV radiation can have on outdoor renderings of objects that use fluorescent colors.

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 appendix.

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 color of the sky typical for such regions. Figure 4.13 shows the subtle but noticeable effects of ground albedo on the sky over a reconstructed archaeological site.

4.7.2 Solar Radiance Model

Figure 4.14 shows two examples of the progression of the direct solar radiance function we obtained for each turbidity level. The color shift that is induced via Rayleigh scattering is clearly visible, as well as the expected reduction in direct solar radiance for decreasing solar elevations.

In figure 4.17, we show close-up views of the sun for the complete model, i.e. both the direct radiance and the sky dome luminance pattern added together. It is noteworthy that the addition of the two components provides a high-quality image of the solar disc with the proper horizontal gradient near the horizon. For higher solar elevations, the model also behaves as expected, but - in accordance with nature - is less visually prominent, in that it "just" provides a very bright solar disc with no discernible gradient on it. For high elevations, the main benefit of the model is that it automatically provides a solar disc of correct brightness to daytime outdoor skies.

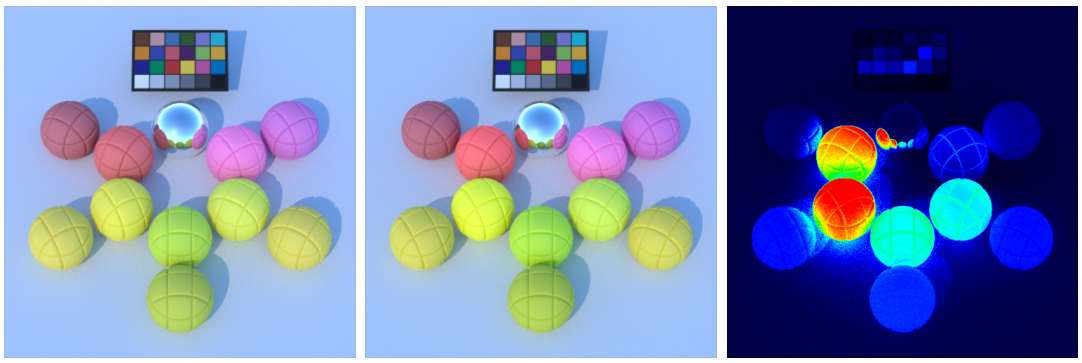

Figure 4.12: The effect of including UV radiation in outdoor renderings that feature fluorescent surfaces. The inner objects have diffuse fluorescent BRDFs with measured re-radiation data of real materials. On the outer objects, only the non-fluorescent main diagonal of the same measurements is used. 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

(a) $\alpha = 0$ (black)

(b) $\alpha = 1$ (white)

(c) 10GY 3/6 (forest green)

(d) 2.5Y 8.5/4 (desert yellow)

(e) 7.5R 4/16 (intense pink)

(f) 1B 5/10 (water blue)

Figure 4.13: 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 color. Reflectance data measured from a glossy Munsell Book of Color. Temple of Kalabsha model courtesy of Veronica Sundstedt.

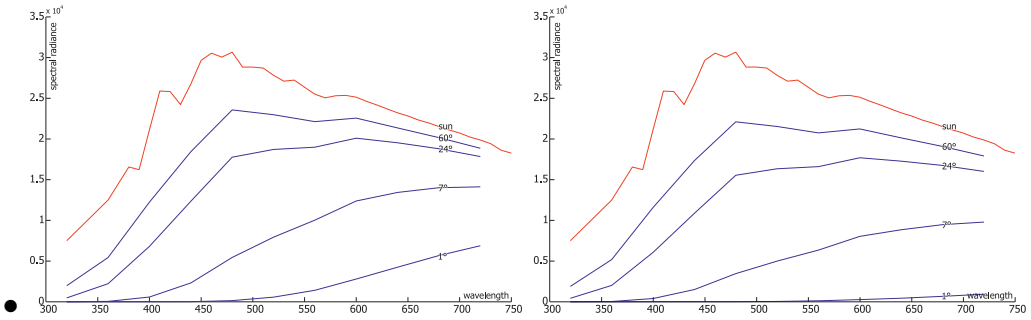

Figure 4.14: Plots of the direct solar radiance for turbidities 1 (left) and 4 (right), and decreasing solar elevations. The spectrum labeled "sun" is the original unmodified solar radiation that hits the top of the atmosphere according to NASA data, as taken from Preetham et al. [1999], with suitable data added for the UV part of the spectrum. Note that the overall tendency of the color shift remains the same, but that the solar radiance for the setting sun in the case of low turbidity is much stronger.

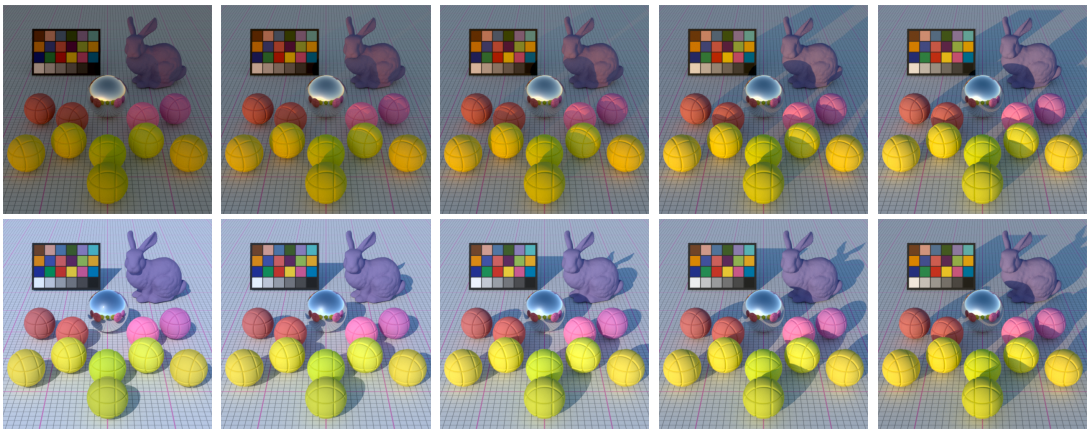

Figure 4.15: Sunrise. solar elevations of (in clockwise order, starting from the top left) 1, 2, 3, 6, 9, 12, 18, 24, 32 and 42 degrees for a turbidity of 3. The images are tone mapped to yield a sequence that approximately follows the brightness development of the sunrise. No white balance adjustment was performed on the images, so that the changing sunlight color can be properly seen.

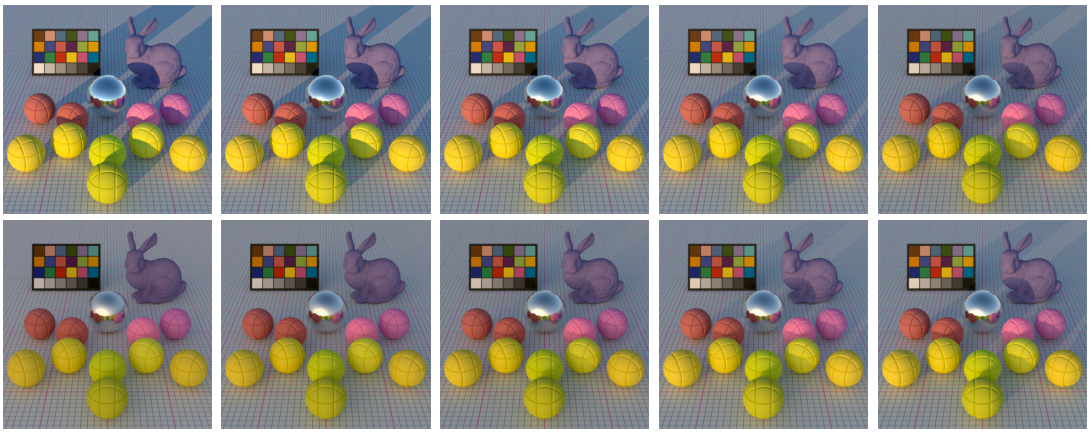

Figure 4.16: **Constant solar elevation, changing turbidity:** for a solar elevation of 8 degrees, turbidity increases from 1 (top left) to 10 in clockwise order. Images are tone mapped to appear roughly equally bright. As with figure 4.15, no white balance computations were done.

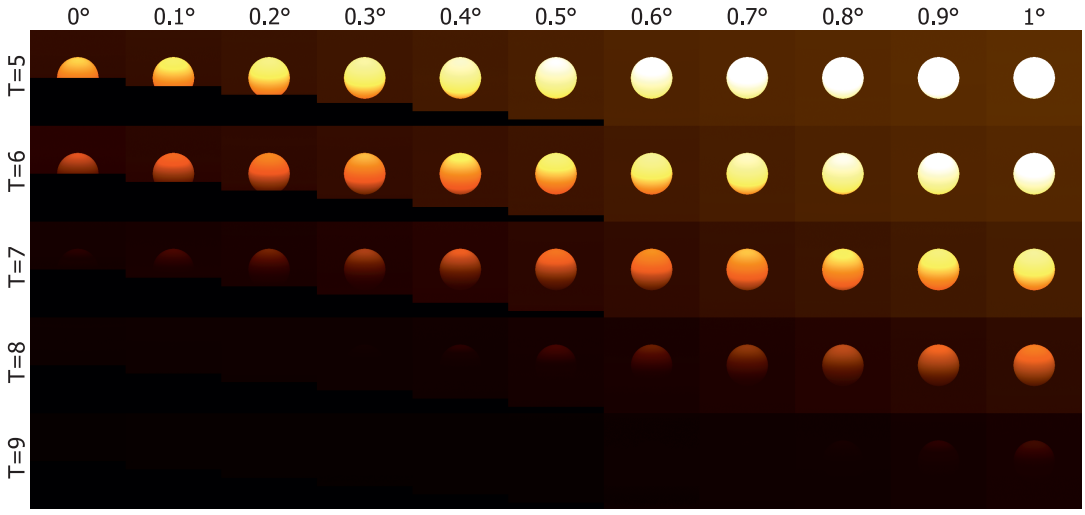

Figure 4.17: Renderings of the region immediately around the rising/setting sun, done with the superposition of the Hošek sky dome model, and the proposed solar radiance function. All small images in the image set are shown at the same exposure level. From left to right, the image sets start with the sun exactly at the horizon, and solar elevation progresses in 0.1° increments upwards. The rows run from turbidity 5 (top) to 9 (bottom). Lower turbidities are omitted, because at this exposure level, they would appear as a uniform white circle. Turbidity 10 is omitted, because the sun would not be visible through the haze.

5. A Full Spectral Analytic Sky-dome model, MK2

In recent years, analytical models of sky-dome and solar radiance have firmly established themselves as useful tools in photo-realistic graphics and visual simulation. With their continuous parameter space of atmospheric conditions and solar elevation, they nicely complement the visually more varied, but also less flexible technique of using HDR captures of real skies to light rendered scenes.

However, even the most advanced analytical skydome models that are currently in use lack some features that would be needed in production environments: in particular, altitude-dependent in-scattering data and haze/absorption data that accurately matches the in-scattered light component for observer altitudes above ground level. Such data is crucial to be able to render arbitrary scenes that include views of distant scenery, such as mountain ranges that recede into the distance, or downward-looking images of aircraft in flight, without having to resort to very time-consuming path tracing simulations. Also, for accurate simulations of specular object appearance in outdoor scenes, it can be important to consider the polarization of skydome radiance – but no such functionality is available for more recent, higher quality models of skydome radiance.

The model described in this chapter is designed to overcome these deficiencies. We follow the general approach of first constructing a brute force simulation (called `atmo_sim`, based on `ART` [Wilkie et al., 2018], a publicly available rendering system that supports spectral rendering and polarization), and then designing an analytical model that can fit the data produced by the reference brute force simulation.

However, we advance the state of the art with regard to practically all components of this process: the underlying atmospheric model now has a more plausible structure and uses a more realistic model of Mie scattering. Due to a modern, unbiased renderer we can run the brute force simulation for a larger parameter space than before: in particular, we now cover below sunset conditions. Furthermore, the model used for fitting has a new functional form that is better capable of capturing luminance patterns. We also extend the model to new capabilities, such as the ability to provide downward-looking data for observer altitudes above the ground level, and a fitted haze attenuation function that matches the atmospheric conditions even for downward-looking directions. And we even provide a model for the polarization state of in-scattered radiance.

Since its original publication in 2012, the model described in chapter 4 was critically reviewed by multiple independent researchers [Kider et al., 2014, Bruneton, 2016]. The reviewers found some flaws in the model regarding its correspondence with real world skydome captures. The new model addresses these issues. Due to more realistic modeling of the atmosphere, in particular the inclusion of an ozone layer and a more realistic vertical profile, particle size distribution and phase function modeling of Mie scatterers, the new model considerably improves on the visual realism of the reference data obtained by brute force simulation. The difference is most pronounced for low solar elevations. In order to handle the now fully spherical radiance function fitting, a new parametric model is introduced,

to replace the family of hemispherical functions introduced by Perez.

5.1 Atmospheric Composition Model

In order to make our experimental set-up as flexible as possible, we define a simple file format called ADF (Atmosphere Description Format) which describes the atmospheric configuration that `atmo_sim` uses when generating images – see appendix A.1 for details on the ADF file format. This way, the atmospheric configuration is completely decoupled from the simulator, and experiments can be run on both synthetic and measured datasets: and as it is a straightforward ASCII format, it is even possible to hand-craft ADF files for specific atmospheric conditions one wishes to simulate.

Given this capability, it is possible to use the atmospheric reference data included with `libradtran` and `OPAC` (see section 2.2.3). For a specific atmospheric type (e.g. continental clean), the `OPAC` package provides a vertical profile and a mixture of aerosol components. Each aerosol component in turn is defined by a particle size distribution and optical properties. This does not map to a scatterer record in an ADF file 1:1 because a single record in an ADF file defines a scatterer of a single specific size – this is done so that the tabulated phase function can be used for importance sampling without any further computationally intensive evaluation. Therefore when converting an `OPAC` aerosol component to ADF, it has to be expanded into multiple records to produce a mix of particles with the corresponding size distribution.

5.1.1 Synthetic Atmospheric Profiles

Figure 5.1: Comparison of a sky dome measurement from Kider et al. [2014] to results of our model. **Left:** image from the Kider HDR capture dataset for 27-05-13, 09:30. **Center:** result obtained with `atmo_sim` using the synthetic atmosphere profile for turbidity 3.5, solar elevation 41°, ground albedo 0. **Right:** difference plot, using a similar color scheme and scale as that in the Kider paper. Note that the comparison was directly done on the spectral data provided by Kider et al. [2014], and not on their HDR images: the tone-mapped image shown here is included for illustrative purposes only. To better match our result image, it was processed to remove the black surroundings, flipped to match the orientation of the spectral data, and color shifted.

In addition to being able to use real atmospheric profiles, we created a tool that produces synthetic profiles. The model described in chapter 4 uses a simple exponential decay of Mie scatterers with regards to altitude. As shown in figure 5.3, we opted for a different approach: we assume a gradually rising inversion layer at altitudes between 500 m and 800 m, with smooth transition to an exponential fall-off above that. The resulting profile qualitatively matches Lidar profiles of real atmospheres with inversion layers, such as those documented in Gryning and Batchvarova [2012], page 617. The reasoning behind using such an inversion layer is that analytical sky dome models are intended to represent days with totally clear skies. And at least in Europe and other mid-latitude locations, such skies very often coincide with the sort of high pressure weather that facilitates the formation of an inversion boundary. For mid-latitude conditions, it would in fact be rather atypical to have a purely exponential fall-off on a clear day. And as figure 5.1 shows, our proposed inversion layer distribution yields results that are reasonably close to the reference measurements of Kider et al. [2014].

The synthetic vertical profile of Mie scatterers follows this empirically derived ad-hoc function:

$$N(h) = a \cdot e^{h/H_r} + b \cdot logsig(\frac{-h + H_t}{H_s}) \tag{5.1}$$

where:

- $N(h)$ is the number concentration at a given altitude $[m^{-3}]$.
- $H_r = 1500$ m is the scale height.
- $H_t = 500$ m is the cutoff altitude of the inversion layer.
- $H_s = 60$ m is the cutoff steepness of the inversion layer.

The sigmoid term was added based on our observation that on hazy days there is usually a steep drop in haze density that coincides with the upper end of the inversion layer (see fig. 5.2). Based on systematic personal observation of skies done by the author and his supervisor, who are both glider pilots, the 500 m cutoff altitude was selected as a representative inversion altitude for clear spring or autumn days in Central Europe.

The parameters a and b are derived as follows:

- At $T = 1$, both $a = 0$ and $b = 0$, as by definition, at this turbidity the atmosphere should not contain any non-gas constituents
- At $T = 2$, $b = 0$ and a has been obtained by non-linear regression such that the optical density of the resulting atmosphere conforms to eq. 4.1
- At higher turbidities, a remains the same as for $T = 2$, and b has been obtained by non-linear regression such that the resulting optical density again conforms to eq. 4.1.

This way, we have obtained turbidities for all integer values of T lying between 1 and 11. Even though the shape of the vertical profile is a somewhat arbitrary

choice, it qualitatively matches lidar profiles of real atmosphere with inversion layers, such as those documented in Gryning and Batchvarova [2012].

Both the measured and synthetic profiles can be used as a reference input for fitting the analytic model. The synthetic model has a turbidity control and can therefore be used to fit an analytic model with continuous turbidity control. The measured profile can be used to fit a model that represents just one specific atmospheric condition. The synthetic atmosphere profile was verified against the skydome capture data provided by Kider et al. [2014] and shows a high degree of correspondence (see fig. 5.1).

The measured profiles are more complex in that they typically do not define just one Mie constituent with its corresponding vertical profile. E.g. the `OPAC` profiles define an atmosphere consisting of multiple types of Mie scatterers (insoluble, water soluble, soot, sea salt, minerals and sulfate droplets), each with its own vertical profile, particle size distribution and optical properties. Such profiles can also be converted to ADF because the format (and, conversely, `atmo_sim`) support an arbitrary number of scatterer types.

5.1.2 The Ozone Layer

Previous graphics publications [Nishita et al., 1993, Preetham et al., 1999] have, in ample detail, described the composition of typical atmospheres with regard to atmospheric properties that affect sky dome appearance for solar elevations above the horizon. However, the ozone layer, a feature crucial for dusk conditions, has mostly been omitted in graphics sky models so far: to our knowledge, the only graphics researcher to include this feature was Kutz [2013]. The importance of ozone for sky dome appearance has of course long been known in atmospheric research literature [Hulburt, 1953], although a recent publication has called the direct correspondence between twilight sky color and ozone into question [Lee et al., 2011]. In figure 5.4, we show the absorption spectrum of ozone, and representative vertical distributions of ozone in the atmosphere.

Figure 5.5 demonstrates the difference in sky dome appearance when one includes ozone for a high solar elevation, and a sunset situation: the blue hue

Figure 5.2: The upper boundary of the inversion layer as seen from an airplane. These images were taken within a few seconds of each other. Notice the very sharp transition between hazy and completely clean air. Also notice how the forward characteristics of Mie scattering makes the haze light up against the sun. Something that cannot be easily shown in photographs, but which is a noticeable feature in reality is the fact that the transparent "window" in the haze layer, which in the photos is to the left of the plane, actually moves with the aircraft.

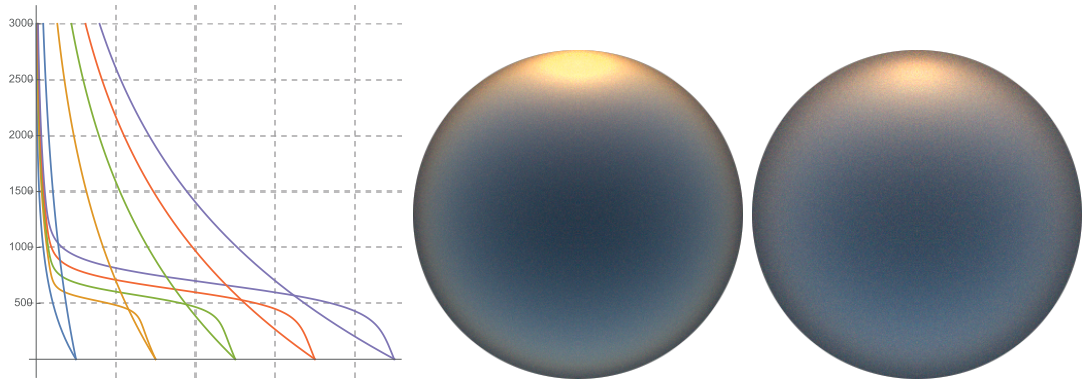

Figure 5.3: The effect of different vertical distributions of Mie scatterers. **Left**: Comparison of scatterer fall off profiles for altitudes up to 3000 m for turbidities 2 (blue), 4 (yellow), 6 (green), 8 (red) and 10 (purple). Matching curves are shown in similar colors. **Center**: turbidity 8, solar elevation 4 degrees, with our proposed inversion layer fall-off. **Right**: turbidity 8, elevation 4 degrees, classical exponential fall-off. Note the significantly hazier appearance caused by the slower exponential fall-off.

of dusk is considerably more realistic if the absorption due to ozone is included. Which makes the inclusion of this feature a necessity, if skies with low solar elevations are to be rendered.

5.2 The Brute Force Simulation

In this section, we give a technical overview of `atmo_sim`, the brute force simulator used to obtain the results shown in this paper. As `atmo_sim` is a path tracer, this section only deals with those properties of the system that are non-standard, such as technical issues that arose in conjunction with polarization support.

5.2.1 The Dataset and Views used for Model Fitting

Side-facing Fish-Eye Views

One significant difference to previous work is that our model now contains the ability to obtain an image of in-scattered radiance for varying observer altitudes. Except at ground level, this means we now have to fit a fully spherical function, and not a hemispherical one, as in previous models. In order to avoid discontinuities at the horizon, we use side-facing fish-eye views of the sky as input to the fitting algorithm, and not two fish-eye views that face up and down. This also allows us to nicely exploit the symmetry of the problem: for the idealized planet defined in the brute force tracer, the spherical in-scattering data is symmetrical around the solar direction, so we only have to generate one of the two side-facing views.

Figure 5.6 shows some sample side-facing views for one particular turbidity and solar elevation: the increasing influence of in-scattered light in the lower half of the hemisphere can nicely be seen, as well as the curvature of the earth that appears for higher observer altitudes.

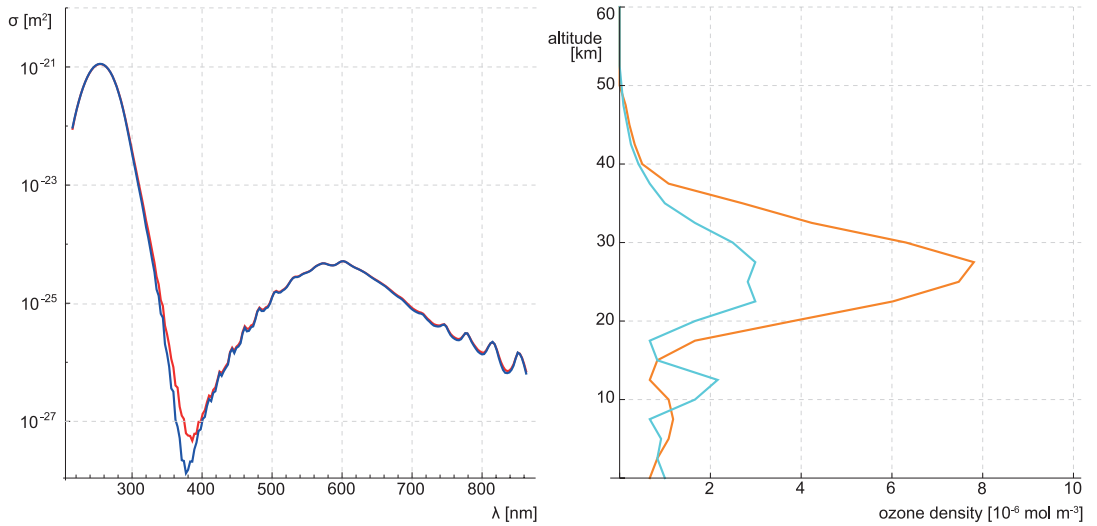

Figure 5.4: The properties of atmospheric ozone. **Right**: the absorption cross section of ozone. The red curve is for 25°C, and the blue one for −50°C, a typical temperature in the ozone layer. We used the −50°C data in our simulation. Data source: Gorshelev et al. [2014]. Note that this plot is on a logarithmic scale. **Left:** a typical vertical distribution profile of ozone in a standard middle latitude atmosphere (orange) and for ozone hole conditions (cyan). Data source: Kaifel et al. [2012].

Parameters

The parameters we used to generate the reference dataset that the fitting process was run on are given in table 5.1. The uneven sampling near the horizon (for solar elevation) and near the ground and the inversion layer (for the observer altitude) was determined iteratively by examination of test renderings: in particular, by examination how fast sky dome features change for varying input parameters in each region of the parameter space.

5.2.2 The Atmospheric Path Tracer

Our brute force simulator, `atmo_sim`, is based on a standard uni-directional path tracer [Kajiya, 1986]. Just like in the previous model, we were faced with the problem that the only source of light in the scene - the sun - covers a very small fraction of the skydome and the probability of hitting it with a ray is tiny. The previous model used a multi-sun optimization, which we did not use here. The reason is that `ART`, the renderer `atmo_sim` is built on, is a general purpose path tracer and we wanted to keep it general, because that allows us to render complex scenes which we use as reference for the transmittance and finite scattering model. Instead we decided to implement a more modern path tracer that uses next event estimation using multiple importance sampling [Veach and Guibas, 1995]. Unlike the previous model, this one does not run on discrete monochromatic wavelengths. When the ray is initially cast, the wavelength is picked randomly from a continuous range and the final value is splatted between neighboring spectral bins.

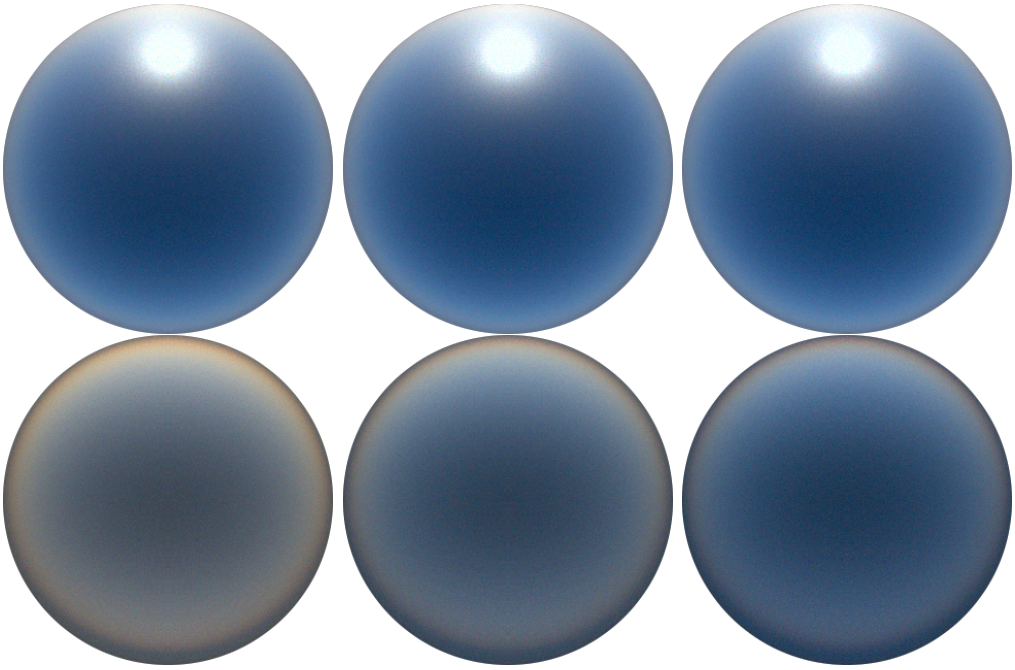

Figure 5.5: The effects of atmospheric ozone on sky dome appearance for turbidity 3. All images in this figure are fish-eye views from ground level. From left to right, the simulated atmospheres contain no ozone, a reduced ozone profile typical for ozone hole conditions (the cyan curve from figure 5.4), and a healthy mid-latitude ozone layer (the orange curve from the same figure). **Top row:** solar elevation 30 degrees. **Bottom**: solar elevation −3 degrees. Note the lack of change for high solar elevations, and the drastically changed hue and luminance of the sky for low solar elevations. Image triplets with similar solar elevation are tone mapped with the same parameters, to correctly show luminance differences.

Scattering Event Sampling

We use Woodcock tracking [Raab et al., 2008, Coleman, 1968] to sample events along path segments. We first sample a distance based on the pre-tabulated minimum mean free path in the ray direction. Then, if the ray neither hits the ground nor leaves the atmosphere within this distance, we sample an interaction event: one of the absorption or scattering interaction types, or "zero scattering". If an absorption event is sampled, the path ends. If Rayleigh or Mie scattering has occurred, the corresponding phase function is sampled and the path continues in a modified direction. If none of these three events are sampled, we say that "zero scattering" has occurred and the path continues unattenuated in the original direction. The probabilities of the event types are determined by their respective scattering / absorption coefficients.

Woodcock tracking is also used to provide unbiased estimates of transmittance along the sun connections in next event estimation. We sample distances along the connecting segment successively until scattering or absorption occurs or the ray leaves the atmosphere. In the former case the transmittance estimate is zero (i.e. the contribution of the connection is zero), in the latter case the transmittance equals one. This procedure is equivalent to the progressive deep shadow maps described by Jarosz et al. [2011].

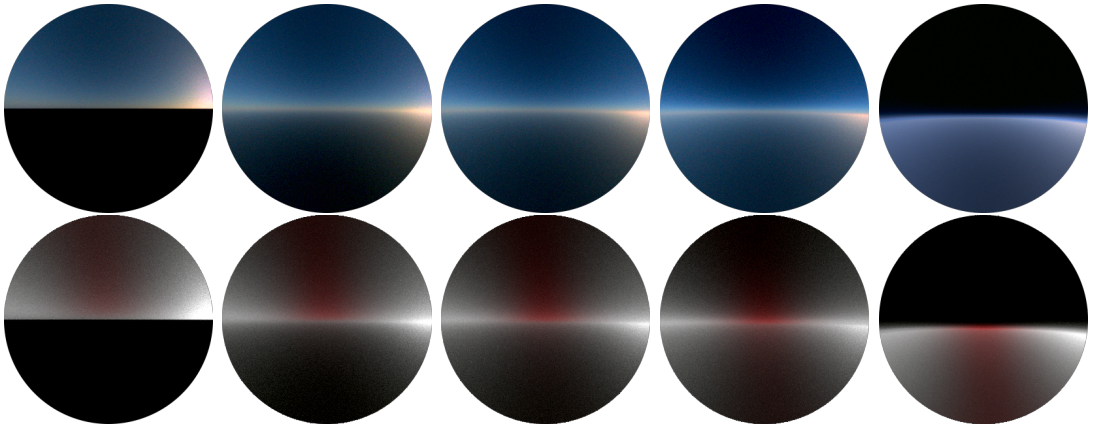

Figure 5.6: Side-facing fish-eye views of in-scattered radiance similar to those used in fitting the model. Turbidity 4, solar elevation 8 degrees, ground albedo 0.3. The solar position is at the right edge of the image. Observer altitudes are (from left to right) 0, 3.5, 7, 15 and 75 km. The latter image clearly shows the curvature of the planet. **Top row**: tone mapped images. Note that the ground is totally black in the left-most images, as only in-scattered light is stored – ground reflections are intentionally omitted. **Bottom row**: scaled overlay of the degree of polarization for 550 nm.

Efficient Ground Albedo Evaluation

To make the dataset generation faster, `atmo_sim` is capable of rendering images for several ground albedos at once. This is possible because the ground albedo affects only the contribution of a path, but not its propagation. So the time-consuming path tracing is performed just once, and the attenuation of Lambert reflection is computed and stored for all considered ground albedos.

Polarization Support

One of the features of our proposed model is the inclusion of skylight polarization. Two main changes have to be made to the otherwise standard path tracing code

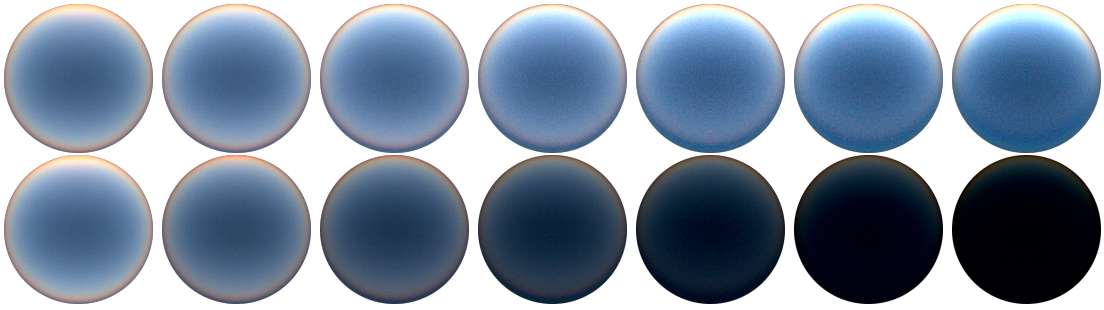

Figure 5.7: Luminance patterns for below sunset conditions. **From left to right**: solar elevations 0 to −6, turbidity 3. **Top row**: images individually tone mapped for similar overall brightness. Note how the shadow of the earth rises opposite the location where the sun went below the horizon, and that the shadow is less and less distinct as night sets in. **Bottom row**: images tone mapped to equal the brightness of solar elevation 0.

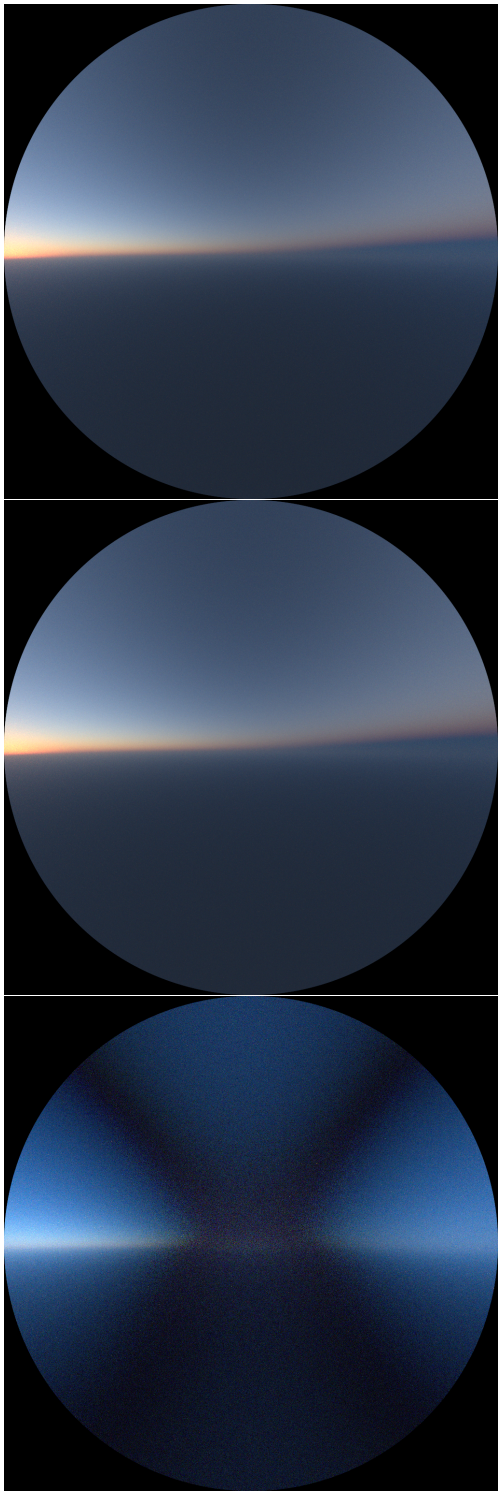

Figure 5.8: Polarization support. **Top:** a reference render of a twilight sky, rendered with polarization enabled. **Middle:** the same scene with polarization rendering disabled. **Bottom:** the difference, multiplied by a factor of 20. The differences are on the border of the perceptible but visible for post-sunset scenarios. For daytime scenarios, the results are virtually identical. Our current theory for this behavior is that post-sunset scenarios are the only ones in which multiple Rayleigh scattering plays a role and such chains of scattering events are liable to be evaluated differently in a polarization-capable renderer.

	Sample values
Turbidity	1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11
Solar elevation (degrees)	-8.05, -5.23, -2.85, -1.00, -0.01, 0.00, 1.00, 2.85, 5.23, 8.05, 11.25, 14.79, 18.64, 22.77, 27.17, 31.82, 36.71, 41.83, 47.16, 52.71, 58.46, 64.40, 70.53, 76.84, 83.34, 90.00
Observer altitude (meters)	0, 1.87, 15.00, 50.62, 120, 234.38, 405, 643.12, 960, 1366.9, 1875, 2495.6, 3240, 4119.4, 5145, 6328.1, 7680, 9211.9, 10935, 12861, 15000
Ground albedo	0, 0.33, 0.66, 1

Table 5.1: Table of parameter values that were used to create the reference dataset used for fitting. Note the uneven sampling of solar elevation and observer altitude, to better sample areas in the parameter space where large changes occur.

to yield such data. First, the image file format used to store the results of the simulator has to support storing complete Stokes vectors for each pixel. Second, a polarization reference frame has to be tracked along with each path, and properly updated when scattering events occur [Wilkie and Weidlich, 2012]. Referring to sections 2.2.1 and 2.2.2, the phase function Mueller matrices are given for reference frames aligned with the scattering plane. The planes of two consecutive scattering events will almost certainly not be the same, which means that the output reference frame and the input reference frame of two consecutive scattering Mueller matrices will also differ. Therefore, before applying a scattering Mueller matrix, the Stokes vector has to be rotated to the correct input reference frame. A Stokes vector S_{out} in a reference frame rotated by ϕ is obtained by applying a rotation Mueller matrix to the original Stokes vector S_{in}:

$$S_{out} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(2\phi) & \sin(2\phi) & 0 \\ 0 & -\sin(2\phi) & \cos(2\phi) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \cdot S_{in} \tag{5.2}$$

This reference frame is also used when computing the attenuation caused by scattering events along the path.

For performance reasons, we also made a third structural change to the path tracer: when working with polarized light, the attenuations that occur along a path should not be concatenated together. For a given wavelength, attenuation samples in a polarization renderer are represented as 4×4 Mueller matrices. When two such samples are concatenated, one of them has to be rotated to the reference frame of the other: this adds significant overhead compared to just rotating the simpler Stokes vector of a light sample. Which means that a polarization-capable path tracer is actually faster if it avoids attenuation concatenation entirely: instead, it retains them separately for the entire path, and individually multiplies each of them with a light sample on its way to the image plane.

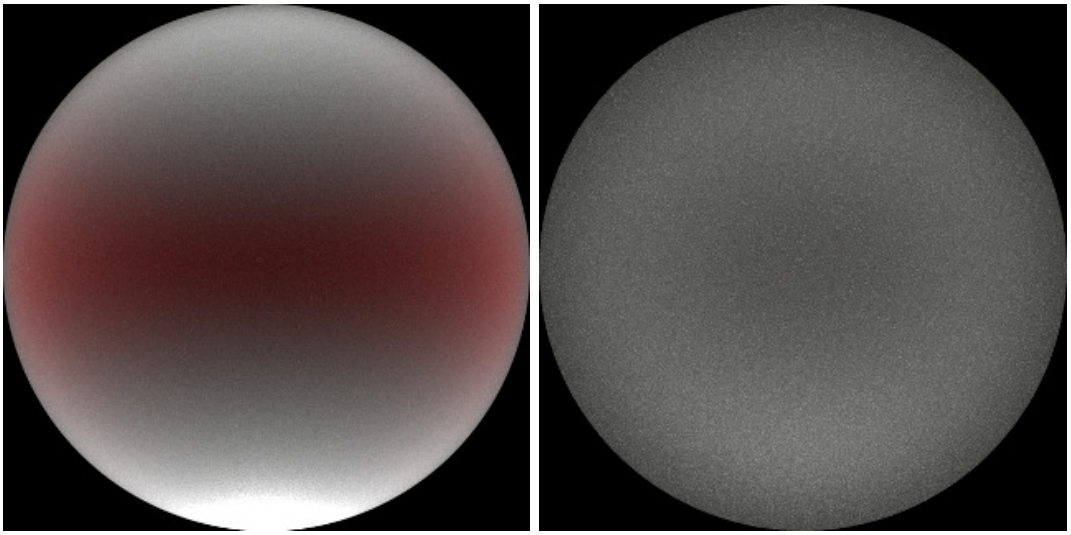

Figure 5.9: Degree of polarization patterns for turbidity 3 (left column) and 11 (right column). **Left:** solar elevation 30 degrees. **Right:** solar elevation 2 degrees. Note the significantly lower degree of polarization for the high turbidity case, resp. the total disappearance of polarization for solar elevation 2.

5.3 The Analytical In-scattered Radiance Model

Initially, we experimented with an approach similar to that of the Perez et al. [1993], Preetham et al. [1999] or our previous model. The approach of designing these models is that the sky-dome radiance pattern consists of certain features – e.g. a bright patch around the sun that falls off exponentially as the solar angle increases. A set of parameters is provided that controls the appearance and strength of these features. These parameters are typically found by the means of nonlinear optimization. This approach works as long as the number of these features is kept low, but that obviously limits the range of radiance patterns the model can reliably reproduce.

We specifically want to be able to capture twilight situations. When the sun is at or below the horizon, two related phenomena can be observed - the Earth casts a shadow onto the atmosphere, producing a wedge of darker color, and above it, a wedge of brighter pinkish back-scattered light (called ”Belt of Venus”, or less romantically ”anti-twilight arch”) is formed. Capturing these patterns would require adding additional terms to a Perez-like model, and with each of these terms the fitting process becomes slower, more memory consuming and prone to getting stuck in local minima. The biggest obstacle however is that it is hard to design the radiance function to consist of orthogonal features – the same radiance pattern can be equally well expressed by two completely different parameter sets, and interpolating between these parameter sets does not produce meaningful results.

Another approach would be to completely forego any attempts at describing the radiance pattern as an analytic function and just store and directly look up the path tracer outputs. The amount of data required to be packed with the model implementation would be huge – our reference images occupy 75 GB. Another option would be to pick points on the sphere, store values for these points and basically draw the sky-dome as a triangulated mesh of these interpolated values.

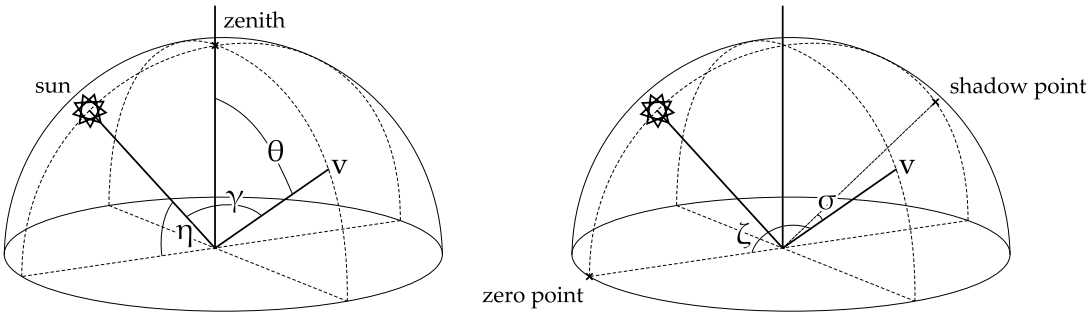

Figure 5.10: Angles used in the analytic model: θ - zenith angle, γ - solar angle, η - solar elevation, σ - shadow angle, ζ - zero angle. v is the view direction.

This approach would have some nice properties – it can be made arbitrarily precise just by increasing the number of control points. Some care would have to be taken when picking the spatial distribution of these points – they would have to be placed densely around the horizon, or in general around the areas with high frequency changes, of which there are plenty: at lower viewer altitudes, there is a steep intensity gradient around the horizon. This happens over the course of 0.5° and the precise point where this happens varies with viewer altitude. Fitting the glow around the sun would also be difficult, and unless the lookup tables were generated in very small solar elevation increments and the placement of control points would be very dense along the sun's trajectory upwards, interpolating between adjacent states would appear as if the sun was fading from one position and at the same time appearing at another position.

We are trying to reconcile two different perspectives: one of them is that a Perez-type model is essentially a phenomenological decomposition of the features visible on the skydome. The other perspective is that such an analytical function is essentially an image compression scheme, highly specialized to compress images of the skydome radiance pattern. We ended up with a tensor decomposition based approach, which, compared to the analytical model described in chapter 4 is a significant shift towards the "compression" paradigm. However before we present the final approach, we briefly describe our previous attempts along with the problems we encountered in hope that the negative results shed some light on the decisions we ended up making in the final model.

5.3.1 The Dead Ends

Feature-Based Phenomenological Decomposition

Let us consider this problem: we have one reference image I produced by the path tracer, representing one specific skydome configuration and conversely one specific skydome radiance pattern, and we are trying to fit it, i.e. find the parameters of our analytic model such that its output approximates the reference rendering.

The model is a function of the view ray direction. We use the same trick as thre Perez model – we represent the view ray direction as a set of angles: the zenith angle θ, the solar angle γ, and a few more angles that we introduce later.

We started with the idea of decomposing the spectral radiance pattern into separate features. The first such feature is the background. The background is a function of one input variable - the zenith angle θ. We can tabulate the function

– store the function values $B(\theta)$ at points $\{\theta_0 = 0°, \theta_1, ..., \theta_n = 180°\}$. The table has to cover values from 0° to 180°, but the sampling points do not necessarily have to be equidistant - the function has a slow rate of change close to 0° and 180°, but there is a sharp change around 90° where the horizon lies. By sampling the function more densely close to the horizon we save space and still get the required resolution in areas that matter.

We constructed the tabulated function the following way: for each sampling point θ_i, the tabulated value $B(\theta_i)$ was obtained by taking the minimum of all pixels in the reference image that lie at angle θ_i relative to the horizon (within a small tolerance). See fig. 5.12, the third column is a visualization of what this background function looks like. The important thing to note here is that the $B(\theta)$ tabulation is not obtained by iterative non-linear fitting but rather by a very simple deterministic approach. We consider that to be a very desirable property because the process is very fast and can not get stuck in any local minimum. When reconstructing intermediate values from the tabulated function, we can use linear interpolation or a more sophisticated approach, such as piece-wise cubic hermite interpolation (PCHIP), which has the nice property of being C^1 continuous.

Now that we have obtained the background function B, we subtract it from the reference image I and we are left with some residuum $R_1 = I - B$ (see the 4th column in fig. 5.12). The residuum image contains some distinct features: the bright area around the sun, the bright areas at horizon directly below the solar point and above the anti-solar point and the Earth's shadow. We try fitting each feature separately – which is why we call this approach feature-based phenomenological decomposition.

Let us take the the bright area around the sun for instance. The obvious approach would be to represent it as a function of one input variable - the solar angle γ. However, we know that the aureole is not completely symmetrical around the sun. We have already noted in our previous model that the glow extends more downwards and horizontally than it does upwards. To fix this, we make the solar glow a function of two variables - the zenith angle θ and solar angle γ, and the function is internally represented as an outer product of two single variable functions, each tabulated:

$$S(\theta, \gamma) = S_{zenith}(\theta) \cdot S_{solar}(\gamma) \tag{5.3}$$

We are now faced with the problem of determining what the tabulated values of S_{zenith} and S_{solar} should look like. Our first naive approach was to start with S_{solar} and define it such that for a given sample at γ_i the value $S_{solar}(\gamma_i)$ would be the maximum of all the pixels in the residuum image that lie at an angle γ_i within a small tolerance. Then the remaining function $S_{zenith}(\theta_j)$ would be the average value of all pixels that lie at an angle γ_j relative to the zenith in an image formed by pixel-wise division of the residuum image by S_{solar}. We hoped that the function S constructed this way is a good enough reconstruction of the solar glow (it is not). We then subtract the solar glow and get s "stage 2" residuum $R_2 = I - B - S$ that contains the remaining features - the front glow, back glow and the shadow.

The remaining features left in R_2 are fitted one by one using the same approach – they are all represented as functions that are outer products of single variable

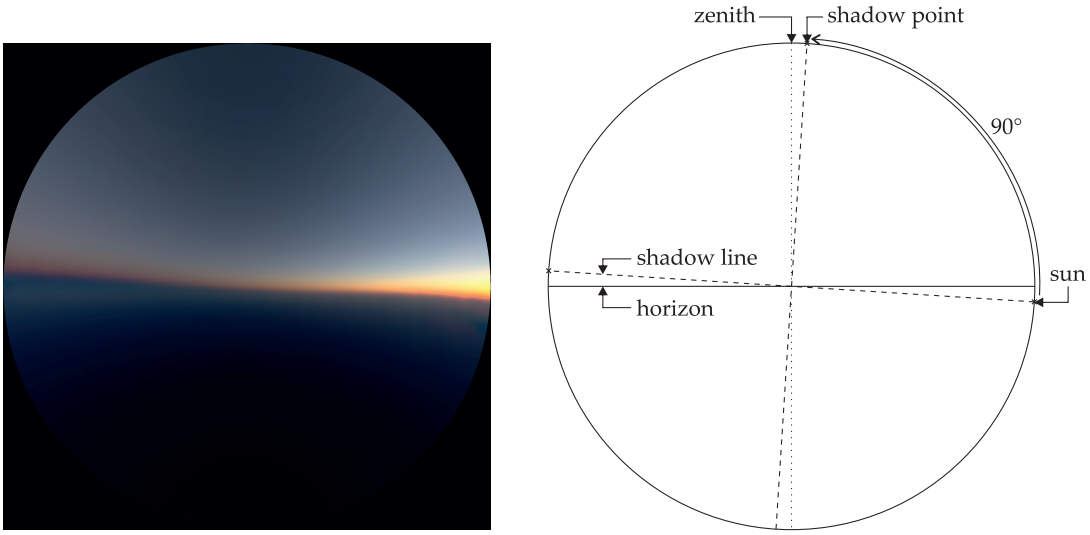

Figure 5.11: A twilight sky, solar elevation $-5°$ **Left:** the reference rendering – a side facing fish eye view. The shadow wedge and the Belt of Venus are visible. **Right:** its geometry. The shadow line forms a ring with a radius of 90° around the shadow point.

functions. The only difference is that since these features are located elsewhere, we have introduced a few more angles: a *zero angle* which is measured relative to a *zero point* placed at the horizon under the sun, and a *shadow angle* measured relative to a *shadow point* placed 90° away from the sun on the direction of the zenith (see fig. 5.10).

There are a few problems with this approach. First, the features are going to interfere with each other. We can sort of work around it by masking out some parts of the image when fitting some feature (e.g. when fitting the solar glow, we can mask out the rear hemisphere to make sure the back glow does not interfere with the fitting). But in some cases the features are impossible to separate: at sunset, the solar glow and the front glow appear at the same location and are indistinguishable from each other, the shadow appears at the same location as the back glow, etc.

Furthermore, the approach of decomposing the solar glow by finding the maxima in one direction and averages in another direction does not produce the best result that we can get by representing a 2D function as an outer product of two 1D functions. In fact, this is exactly the problem that is optimally solved by Singular Value Decomposition (SVD) Low Rank Approximation [Eckart and Young, 1936].

SVD decomposes a matrix (in our case the 2D tabulated function) into a sum of outer products of vectors (our 1D functions). If we use just a rank 1 approximation, that would essentially constitute our 1D function tables. and it would be optimal in the sense that no other outer product of two vectors can approximate the function better.

But again, there are a few problems with this approach: first, some of the elements of the function to be decomposed are undefined because not all combinations of zenith and solar angle correspond to an actual direction (e.g. there is no direction corresponding to $\theta = 0°$ and $\gamma = 0°$ unless the sun is at the zenith).

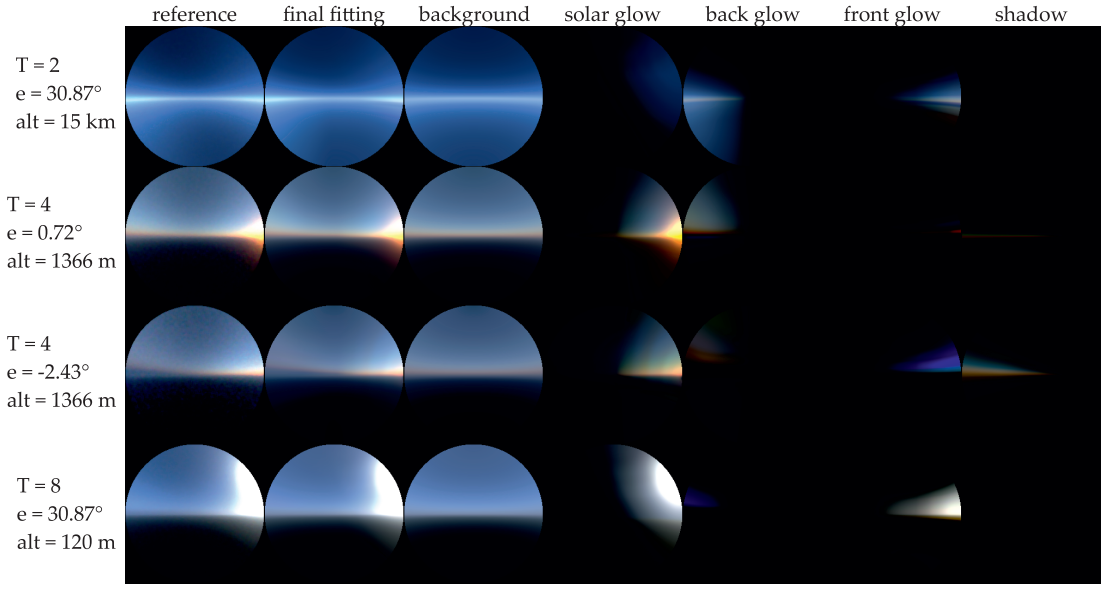

Figure 5.12: Radiance pattern, feature-based fitting attempt. 180° side fish-eye views of the sky and in-scattered light.

Second, and most importantly, it does not solve the issue of features interfering with each other.

Tensor-Based Decomposition

The next logical step is to forego any attempt at explicitly separating the features. Instead, we let the algorithm handle the optimal decomposition completely. Our previous attempt used four input angles angles (sun, zenith, zero, shadow), so we just attempted to replace the whole model by a sum of outer products of four variables. The whole function that computes the spectral radiance for a given ray direction can then be calculated as:

$$\mathbb{F}(\theta,\gamma,\sigma,\zeta) = \sum_{i=1}^{n} \mathbb{F}^{(i)}_{zenith}(\theta) \cdot \mathbb{F}^{(i)}_{sun}(\gamma) \cdot \mathbb{F}^{(i)}_{shadow}(\sigma) \cdot \mathbb{F}^{(i)}_{zero}(\zeta) \tag{5.4}$$

We now start with a tabulated 4D function $\mathbf{F}$ which we decompose into a sum of n sets of four outer products. This approach has a some indisputable advantages: the decomposition does not require that we do any explicit separation of features, but most importantly, *we can approximate the radiance pattern to arbitrary precision.* If more precision is required, we can simply increase the sampling density (i.e. the dimensions of the tabulated function $\mathbb{F}$) or the decomposition rank n.

The function tabulation is a 4D array, so we use a tensor generalization of SVD called Canonical Polyadic Decomposition (CPD) [Kolda and Bader, 2009]. The algorithm produces an optimal decomposition analogue to SVD low rank approximation. It can also deal with tensors that contain undefined elements, so in theory all our problems should be solved.

However the problem of undefined elements needs to be examined more closely. We start with a 2D image, which we expand into the 4D tensor. The four dimensions of the tensor correspond to the $(\theta,\gamma,\sigma,\zeta)$ input variables. Not all of these

angle combinations will be valid – in fact, since we started with a 2D image, the set of valid angle combinations forms just a curved 2D surface in the 4D tensor space and as a result, a vast majority of the tensor is undefined, which makes the decomposition algorithm numerically unstable and also wastes space.

5.3.2 The Final Model

The whole approach of the tensor-based decomposition is still sound, the only problem is that we started with the idea of "features," which led us to the impasse where we tried to add more angles centered around each feature in hope that it would make fitting the radiance pattern easier or more efficient. In the final model, we scale back a bit and return to a function of two (albeit differently selected) variables. Decomposing a 2D tabulated function makes the fitting stable and efficient. Thus we have come full circle and returned to storing the whole image essentially just compressed using SVD low rank approximation. There are, however, still a few issues that need to be addressed:

The Input Space

SVD decomposes the input into a sum of separable matrices. A separable matrix is a matrix that can be expressed as an outer product of two vectors. Separable matrices are, in a manner of speaking, axis aligned. Transforming the input image into a space where the features of the image are axis aligned makes the low rank approximation work better.

The first input parameter is the solar angle γ. This makes the gradient of the solar glow roughly parallel to the first axis – *roughly* because, as we have demonstrated earlier, the solar glow does not form a perfect circle around the sun, but tends to extend further to the sides and downward.

In addition to the features of Preetham-like models, we now specifically also want the model to capture Earth's shadow well. Therefore when fitting images of post-sunset skies, we have chosen the shadow angle σ to be one of the input parameters that defines one of the principal axes. See fig. 5.11 for reference – the shadow line is perpendicular to the solar angle gradient and parallel to the shadow angle gradient, which makes it perfectly aligned with both axes.

For regular daytime skies, no shadow is visible and using the shadow angle as an input parameter provides no benefit since its axis is not parallel or perpendicular to any feature visible on the sky. We instead use the zenith angle θ as the second parameter. The axis defined by the zenith angle is perpendicular to the horizon, which makes it easier to accurately represent it with a low rank

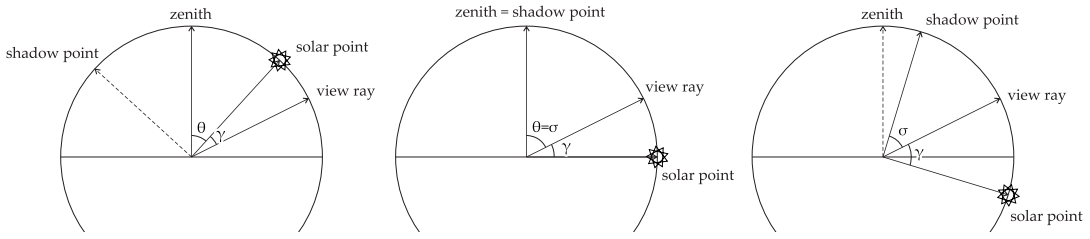

Figure 5.13: The input angles. Left to right: daytime, sunset, after sunset.

approximation.

The model switches between these two modes at solar elevation 0° when the zenith point aligns exactly with the shadow point, which makes the transition completely seamless (see fig. 5.13). To present this approach in a unified manner, we introduce an angle α, which we for lack of a better name call simply the *zenith/shadow angle*, which is equivalent to the zenith angle for solar elevations greater than 0° and equivalent to the shadow angle otherwise. Note that this makes the horizon not aligned with any axis in post-sunset skies. The image emphasis process described later was developed to make fitting of the horizon line more accurate in these conditions.

We end up with a function of two parameters γ, α, which is internally represented as an outer product of two tabulated single parameter functions:

$$\mathbb{F}(\gamma, \alpha) = \sum_{i=1}^{n} \mathbb{F}^{(i)}_{solar}(\gamma) \cdot \mathbb{F}^{(i)}_{zenith/shadow}(\alpha) \tag{5.5}$$

The reference image, which is produced by the path-tracer as a fisheye projection, has to be transformed to the (γ, α) space (see fig. 5.14). The model fitting procedure is then the procedure of finding the tabulated functions $\mathbb{F}_{solar}$ and $\mathbb{F}_{zenith/shadow}$.

We henceforth refer to the image remapped to the (γ, α) space as the *tensor*. We have chosen the dimensions of the tensor to be 361×361 (i.e. 0.5°/pixel since the valid values of γ and α are 0° - 180°). The tensor is then decomposed into an outer vector product using the CPD low rank approximation algorithm. Even though the tensor is just two-dimensional and could be processed by a standard SVD, we still use CPD because it natively deals with undefined values and as figure 5.14 shows, not all combinations of angles are valid.

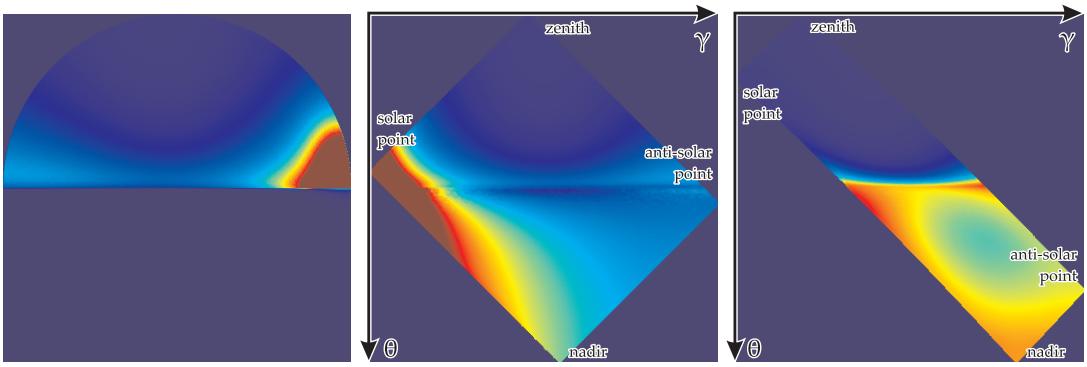

Figure 5.14: **Left:** The input fisheye image. Single spectral channel in false color. Solar elevation 8°. Altitude 15 m above sea level, hence the image is almost completely zero below the horizon line. **Center:** The same image transformed to the (γ, α) space. The valid combinations of γ and α form a parallelogram. The below-horizon part of the image is amplified by the pre-emphasis process. **Right:** The (γ, α) remapping of a different skydome, solar elevation 55°, which is why the valid angle combinations form a different shape.

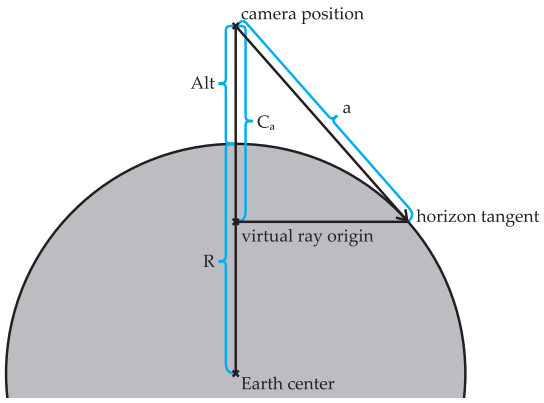

Figure 5.15: Horizon correction geometry. R - Earth radius, Alt - camera altitude, C_a - correction length, a - tangent length.

High Altitude Angle Correction

The analytical model is shipped with a finite number of fitted skydome configurations. Intermediate states have to be interpolated. For instance, we might only have the fitted data for solar elevations $\eta = 20°$ and $\eta = 40°$ and we might want to render the sky at a 30° elevation. Or we might only have the data for altitudes 5 km and 15 km and we might want to render the sky at a 10 km altitude.

It is worth noting that if we take the fitted data for a specific solar elevation (e.g. $\eta = 20°$) and use them to draw an image of a different elevation (e.g. $\eta = 40°$), the *direction* $\leftrightarrow (\gamma, \alpha)$ mapping process works in our favor and automatically moves the solar point and the aureole around it at the right position. The resulting brightness pattern will obviously not correctly reflect the change of the sun's position, but the geometry will be in the right place. Therefore, if we need to render a skydome for $\eta = 30°$, but we only have the fitting data for $\eta = 20°$ and $\eta = 40°$, we can render two images using these two provided adjacent fittings treated as $\eta = 30°$ fittings, and interpolate between them. The interpolation will not produce two half transparent suns at 20° and 40°, but rather a single sun at 30° with an appearance that is a mix of the two provided fittings.

The same can unfortunately not be said for the horizon line. The horizon line only appears at $\theta = 90°$ for altitude = 0. For all altitudes higher than 0, the horizon line appears slightly below the observer, and this effect is more apparent at higher altitudes. If we had tabulated data for just two altitudes, e.g. 100 m, and 100 km, an animated transition between these states would appear as if the horizon line was disappearing from one position and appearing at another instead of the desired effect of the horizon line gradually moving from the initial position to the final position. The reason is, view rays that are tangent to the Earth's surface are at different θ for different camera altitudes (see fig. 5.15). In real scenarios we would obviously not interpolate between two states that are that far apart, but this problem still has to be addressed if we want to render animations with smoothly changing camera altitude.

To solve this issue, we propose a correction to the way the θ and σ angles (and thereby also α) are calculated. The principle of this correction is to ensure that if the view ray is tangent to the Earth's surface, its θ will always be 90°, regardless of the camera altitude.

The un-corrected way of calculating the angles is as follows: the view direction is given as a unit vector $\overrightarrow{v}$. Likewise, the directions towards the solar, zenith and shadow points are given as unit vectors $\overrightarrow{s}$, $\overrightarrow{z}$ and $\overrightarrow{u}$ (we use u as *umbra* because s is already taken). Assuming that the z-axis points upwards, $\overrightarrow{z} = (0, 0, 1)$. The angles are then calculated simply as:

$$\begin{aligned} \gamma &= \cos^{-1}(\overrightarrow{v} \cdot \overrightarrow{s}) \\ \theta &= \cos^{-1}(\overrightarrow{v} \cdot \overrightarrow{z}) \\ \sigma &= \cos^{-1}(\overrightarrow{v} \cdot \overrightarrow{u}) \end{aligned} \tag{5.6}$$

To correct angles θ and σ, we define a point called *virtual ray origin*, which is a point lying on a line between the camera position and the Earth center, positioned such that if we cast a tangent to the Earth's surface from the real camera position, the point where it touches the Earth's surface lies at a $\theta = 90°$ direction from the virtual ray origin. See fig. 5.15 for the details. Euclid's triangle formula tells us that

$$\begin{aligned} C_a &= R + Alt - \frac{R^2}{R + Alt} \\ a &= \sqrt{(R + Alt)^2 - R^2} \\ \overrightarrow{corr} &= (0, 0, \frac{C_a}{a}) \end{aligned} \tag{5.7}$$

$\overrightarrow{corr}$ is the correction vector. The corrected view direction is then calculated as

$$\overrightarrow{v'} = \text{normalize}(\overrightarrow{v} - \overrightarrow{corr}) \tag{5.8}$$

This corrected view direction $\overrightarrow{v'}$ is then used to calculate the zenith and shadow angles (not the solar angle!).

Improving Horizon Fitting With Image Emphasis

Now that we have defined the input space in the above mentioned way, we are facing the problem that for post-sunset configurations, the horizon is still there and it is not aligned to any axis. Recall that the α angle is defined in such a way that the horizon forms a horizontal line in the tensor only for daytime solar elevations. To fix the issue, we have developed a method that removes the sharp horizon from the image. This process, which we have called *pre-emphasis* is performed prior to transforming the reference image into a tensor and fitting that tensor. At run time when using the model, the output of the fitting is then transformed into its original state using an inverted process called *de-emphasis*, which returns the sharp horizon transition back.

The pre-emphasis process works as follows: We denote I_{orig} the original input image and $I_{orig}(x)$ its value at pixel position x. The input image I_{orig} is cut into two parts very close to the horizon line (we chose a value of $30'$ above the horizon). The top/above horizon part is left intact. The bottom / below horizon part is deleted and infilled using Matlab's standard `regionfill` algorithm that smoothly interpolates inward and solves the Dirichlet boundary condition. This

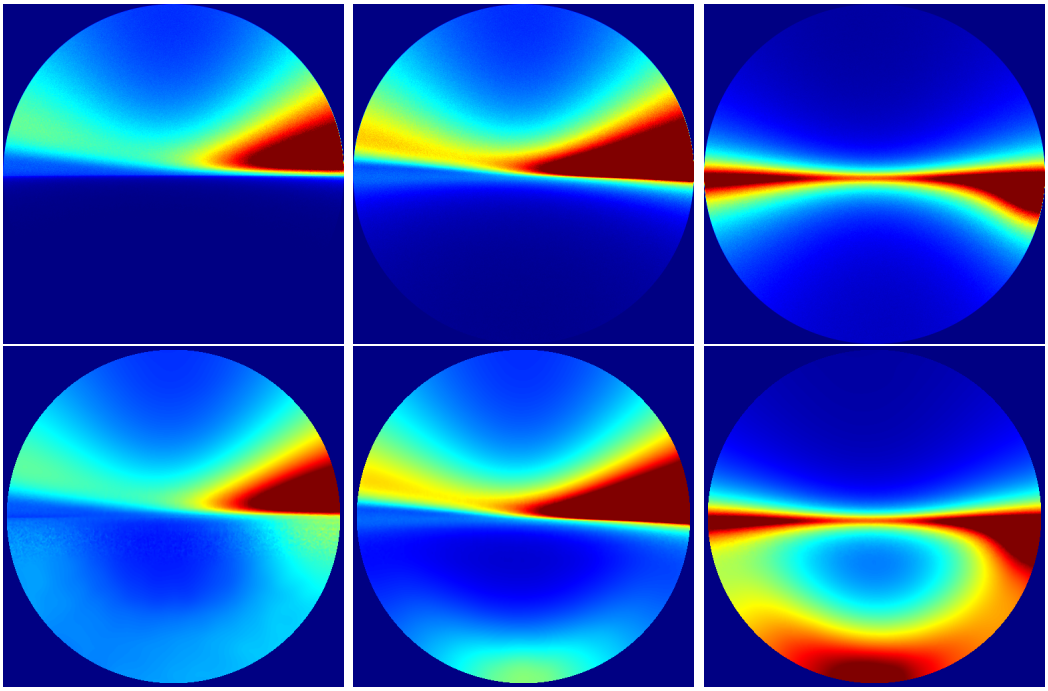

Figure 5.16: The pre-emphasis process. **Top row:** the original fisheye images, prior to any processing, in false color. **Bottom row:** the same images after pre-emphasis has been applied. **Left column:** the g therm fully kicked in, bottom part of the image is completely replaced with in-filled data. **Center row:** g term partially kicked in, making the whole bottom part of the image slightly brighter. **Right column:** only the θ term active, the in-fill is gradually blended in by angle.

essentially removes the sharp horizon transition and leaves the bottom part of the image a completely featureless smooth gradient with brightness equivalent to that of the skydome area just above the horizon. We denote this newly in-filled image as I_{infill}. We also note a single value g denoting the ratio of the average brightness of the original below-horizon area relative to the brightness of the newly in-filled area.

By linearly interpolating between I_{orig} and I_{infill}, we create a guide image I_{guide}. The control value of the linear interpolation is the ratio g. When g is high, meaning that the below-horizon part of the image was already bright enough in the original, we mostly leave the original intact, except for the very bottom part near the nadir, which we always replace by the infill, because it's generally always noisy.

$$\begin{aligned} I_{guide}(x) &= \text{lerp}\left(I_{orig}(x), I_{infill}(x), V_c\right) \\ V_c &= \text{sat}\left(\text{sat}\left(\frac{\theta(x) - \pi/2}{\pi/2}\right) + \text{sat}\left(\frac{g - 0.5}{-0.2}\right)\right) \end{aligned} \tag{5.9}$$

Where:

- V_c is the control value of the linear interpolation.
- sat is the *saturate* function that clamps the value to $[0, 1]$.

- $\theta(x)$ is the zenith angle value at pixel location x.

The control value V_c consists of two terms: the first one makes sure that the original image is always gradually replaced by the in-filled image, starting at $\theta = 90°$ (at the horizon) and progressing towards $\theta = 180°$ (the nadir). The second term makes sure that if the bottom part of the original image was too dark, it is replaced completely, the effect being gradually applied depending on g, starting at $g = 0.5$ and being finishing at full strength at $g = 0.3$.

Having the guide image ready, we calculate a pixel-wise ratio between the guide and the original image:

$$I_{ratio}(x) = \frac{I_{orig}(x)}{I_{guide}(x)} \tag{5.10}$$

The ratio image is then used to define the de-emphasis, a function of zenith angle θ:

$$E(z) = \underset{x;\ \theta(x)=z}{\text{mean}} I_{ratio}(x) \tag{5.11}$$

The function $E(z)$ is tabulated and becomes part of the model data. Next, the pre-emphasized image is calculated by applying the inverse of the de-emphasis function:

$$I_{preemph}(x) = \frac{I_{orig}(x)}{E(\theta(x))} \tag{5.12}$$

This pre-emphasized image is then used as an input for the tensor-based decomposition.

The Fitting Process

We now have all the components required to perform the fitting. To recapitulate the whole fitting process:

1. We start with I_{orig}. This is the reference rendering of the skydome produced by the path-tracer.
2. The pre-emphasis process is performed on I_{orig}, producing a pre-emphasized image $I_{preemph}$ and a tabulated de-emphasis function E.
3. $I_{preemph}$ is transformed into the tensor, essentially producing a two-dimensional tabulated function $\mathbb{F}(\gamma, \alpha)$.
4. The tensor is partially in-filled and filtered. More on that later in this section.
5. A CPD decomposition is performed, producing a set of pairs of one-dimensional tabulated functions $\mathbb{F}^{(i)}_{solar}$ and $\mathbb{F}^{(i)}_{zenith/shadow}$.

The final products of the fitting are:

- The de-emphasis function E
- The tabulated functions $\mathbb{F}^{(i)}_{solar}$ and $\mathbb{F}^{(i)}_{zenith/shadow}$.

These constitute all the data required to render the skydome using the analytical model.

The last step that needs explaining is the in-filling and filtering of the tensor. In the section about high altitude angle correction, we made a quite bold claim that the data produced by fitting one skydome configuration can be used to to render a different skydome configuration. And yet, if we look at the tensors in fig. 5.14, we can see that two different solar elevation configurations produce tensors with valid areas that do not quite overlap. Fig. 5.24 shows what happens when we we attempt to use a fitting of a specific solar elevation to draw an image of a different solar elevation – the rainbow artifact is caused by the fact that the model reads from parts of the tensor that are undefined and basically contain just garbage.

To fix this, the tensor has to be partially infilled. The procedure is really simple: The valid area or the tensor is dilated and the missing data is filled using Matlab's `regionfill` algorithm. Deciding how large an area of the tensor should be in-filled is a trade-off: in-filling a larger area will allow the fitted data to be used to draw skydome configurations that are further away from the original configurations. However fitting a tensor with a larger valid area comes at the price of decreased accuracy in its original valid area. Therefore, we infill each tensor just enough to cover the intermediate values between the provided fittings (see tab. 5.1).

Figure 5.25 shows how the in-filling fixes the problem and enables us to perform this solar angle reprojection. Figures 5.27, 5.28 and 5.29 showcase another type of reprojection – this time drawing an image at a different observer altitude – and how the in-filling fixes the artifacts.

After the tensor is in-filled, it has to be filtered. The CPD decomposition is very sensitive to noise - noise on the input produces noise on the output. However, the output noise is not in the image space, but in the (γ, α) space – instead of a grain-like effect, it produces ringing artifacts. To get rid of as much noise as possible, the reference renderings are filtered (using Matlab's `wdenoise` wavelet noise reduction algorithm), and after the image is reprojected to a tensor and in-filled, it is again filtered, this time using a gaussian blur with $\sigma = 2$ in the solar axis and $\sigma = 1$ in the zenith/shadow axis (recall that the resolution of the tensor is 0.5° per pixel). Fig. 5.31 showcases the problem and fig. 5.30 shows how a proper filtering fixes the artifact.

Due to the nature of the images we are trying to produce – large areas of almost uniform color with very smooth gradients – the final fitting has to be completely devoid of any such ringing artifacts because even a slightest unevenness is *very* apparent on the final renderings.

Regarding the decomposition itself, we have chosen to use decomposition rank $n = 5$, i.e. the decomposition produces five sets of vectors. In all images we have tested, $n = 5$ produces a decomposition that explains > 99.5 % of the variance in the tensor. Choosing higher ranks produces better results on paper, e.g. $n = 10$ produces a decomposition that explains > 99.999 % in all cases, but the results are not any different visually, it would just make the dataset $2\times$ bigger.

The tensor decomposition produces vectors of length 361 (0.5° increments). These do not have to be distributed in their entirety in the final model: The essentially represent tabulated functions and the samples do not have to be placed

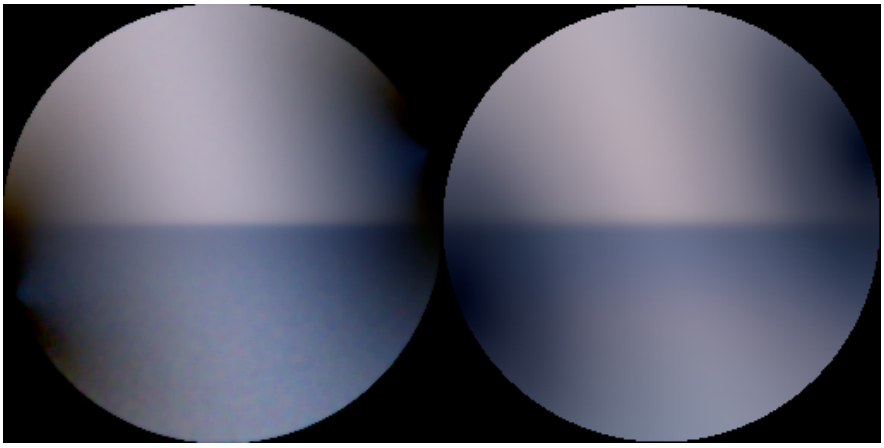

Figure 5.17: The raw polarization data of the reference output and its corresponding fit. Turbidity 6, solar elevation 19.44°, altitude 2495 m.

uniformly: we have chosen to sample $\mathbb{F}_{solar}$ densely at lower angles (areas directly surrounding the sun) and sparsely around the anti-solar point, giving us satisfactory results at 44 samples. Similarly, $\mathbb{F}_{zenith/shadow}$ is sampled sparsely around the zenith and nadir and densely around the horizon, giving us 41 samples. The same approach has been used for the de-emphasis function, which is sampled more densely around the horizon, giving us 45 samples.

5.3.3 Using the Analytic Model

The user of the analytic model is provided with the model implementation as well as the fitted data. The data is provided as the tabulated functions $E(\theta)$, $\mathbb{F}^{(i)}_{solar}(\gamma)$ and $\mathbb{F}^{(i)}_{zenith/shadow}(\alpha)$. For a given view direction $\overrightarrow{v}$ which the user wants to evaluate, he calculates the angles $\gamma, theta, \alpha$. The spectral radiance coming from that direction is then calculated as:

$$L = E(\theta) \cdot \sum_{i=1}^{n} \mathbb{F}^{(i)}_{solar}(\gamma) \cdot \mathbb{F}^{(i)}_{zenith/shadow}(\alpha) \tag{5.13}$$

The fittings are provided for a finite number of elevation / altitude / turbidity / albedo / wavelength configurations (see table 5.1). When a configuration that is not exactly one of the provided combinations is desired, it has to be interpolated from adjacent configurations using the geometry reprojection approach outlined in the previous section.

Evaluating the model is computationally very simple and can be implemented as a simple shader program – implementing formula 5.13 is just 11 texture fetches, 6 multiplications and 5 additions.

5.3.4 Accompanying Models

The Polarization Component

Building an unoptimized analytical model for sky dome polarization would require to fit functions to the three additional Stokes components present in each polarized wavelength sample. However, due to the nature of Rayleigh and Mie

scattering, sky dome polarization is almost entirely linear, which reduces the problem to fitting the second and third components: together, these encode the entire linear polarization component in a Stokes vector. And for sky domes, one can take this even further: as illustrated in figure 5.18, one can simplify the problem by suitably rotating the reference frame of all the pixels in the fish-eye images used for the fitting process.

As result of this re-parameterization, one can resort to only fitting the second Stokes component: the polarization information encoded in this modified channel has to be rotated back before use, but otherwise is a good representation of the dominant features of the phenomenon. This idea of such a polarization pixel coordinate system re-alignment is not new: Wilkie et al. [2004] already used a similar approach for their model. However, as their technique was only based on indirect reasoning, they had no solid evidence that this sort of re-parametrization is actually permissible. As we ran the entire brute force simulation with polarization enabled, we were able to actually assess the impact removing the third Stokes component has. While it does introduce an error insofar as the polarization effects of higher order scattering are discarded, figure 5.18 shows that after re-parameterization, the third channel only contains a very weak signal that has no perceptible effect, even in scenes that exhibit differences due to sky-dome polarization. The third component signal proved to be similarly weak for all result images of the fitting dataset, so omitting it was a safe choice.

The approach to fitting the polarization pattern is the same, however as the

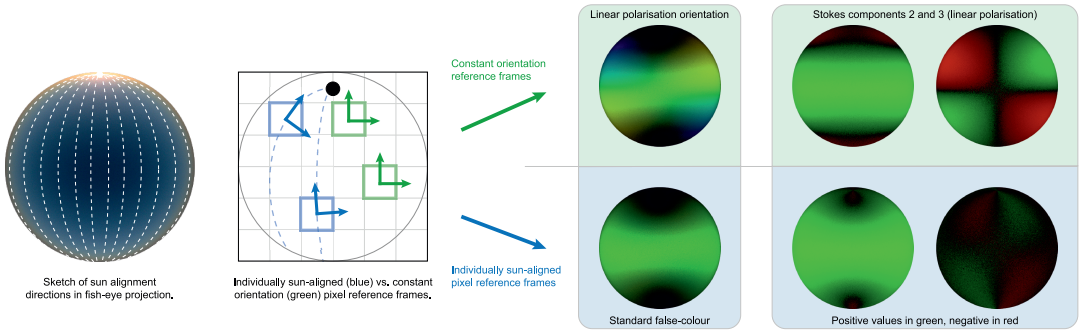

Figure 5.18: Illustration of the sun-aligned polarization reference frames that allow us to only work with the second Stokes Component of polarized sky dome images. Standard reference frames are shown in green, individually aligned ones in blue. The reference frames are right-hand coordinate systems, as the light is coming from the image plane towards the observer. Note that even with the special alignment, the third Stokes Component is not zero, but just very weak: this component is caused by multiple polarizing scattering events, which are rare (see fig. 5.8). We found that for all turbidities, solar elevations and observer altitudes used in the model, the sun-aligned third component is so small that it can be omitted for rendering purposes. For this example, a sky dome with turbidity 3 and solar elevation of 5 degrees was used. Note that in this figure we demonstrate the concept with upward-facing fish-eye views, as these give a view of the complete polarization pattern. For our actual model fitting, we used side-facing fish-eye views, as discussed in the text: the re-orientation works the same in either case.

pattern is much simpler we can make reduce the approximation rank to $n = 2$ and also reduce the sizes of the individual function tables.

Atmospheric Transmittance

A method for fitting the atmospheric transmittance function that matches the skydome model has been developed by Thomas Bashford-Rodgers. It is included in appendix A.2 since it is an integral part of the complete model and will be published together.

5.4 Discussion and Results

The whole model with the whole settings range that we want to support would contain all combinations of 11 spectral bands, 11 turbidities, 21 altitudes, 26 solar elevations and 2 albedos, giving a total of 132,132 combinations. Producing the whole dataset is a work in progress, with the bottleneck being rendering of reference images for all the configurations. We are investigating possible optimization approaches which include the hero wavelength sampling approach [Wilkie et al., 2014] and the zero scattering optimization approaches proposed by Miller et al. [2019]. We have, however, rendered and fitted a representative cross-section that contains both low and high altitudes, turbidities and solar elevations and confirmed that the analytical model can fit these reference images with sufficient accuracy (in fact, with arbitrary accuracy if we produce a sufficiently high rank approximation, but $n = 5$ seems to be enough for all tested cases). The fitting process itself is fast, with one configuration taking less than one minute on a 6-core Haswell-E machine.

As mentioned before, we determined that using approximation rank $n = 5$ produces sufficient result. We have used these numbers of records in the function tables:

- 45 records for a pre-emphasis function $E(\theta)$ table
- 44 records for a $\mathbb{F}^{(i)}_{solar}(\gamma)$ function table
- 41 records for a $\mathbb{F}^{(i)}_{zenith/shadow}(\alpha)$ function table

giving us a total of $45 + (44 + 41) \times 5$ records per a single configuration. With function tables of this size, the SNR error measure never dropped below 35 dB in our representative sample of skydome configurations, which is a significant improvement over the previous model which in some configurations produced SNR as low as 8 dB. It should also be noted that the worst error measurements were obtained at low solar elevations where the noisiness of the reference image already played a significant role.

The uncompressed size of the full fitted model data will be just under 60 MB.

5.4.1 Comparison to Existing Models

The newly presented model brings significant advancements over pre-existing skydome models [Preetham et al., 1999, Hošek and Wilkie, 2012]: it provides data

for downwards facing view directions and solar elevations below 0°, comes with a transmissivity model, which in conjunction with the sky-dome component can be used to evaluate in-scattered light at finite distances, provides the ability to set observer altitude, and includes a polarization component – the only existing model of sky-dome polarization [Wilkie et al., 2004] is known to be approximate at best. With the correct input parameters, it can be set to match the reference data, as provided by the Kider et al. [2014] paper.

5.4.2 Aerial Perspective

Figure 5.19 shows results for the combination of directly viewed sky dome radiance with attenuation and in-scattered light for finite viewing distances, as discussed in section A.2.1. In such a setting, the fitted model is several orders of magnitude faster than the brute force path tracing of the atmospheric model it replaces, which required tens of thousands of samples for convergence. For the mono-wavelength spectral MIS path tracer we integrated the fitted model into, directly viewed diffuse surfaces are basically noise-free after approximately 512-1024 samples per pixel. Which is a fairly typical value for such path tracers and complex environment lighting (a hero wavelength spectral renderer would be faster), so the sky dome model itself is not a constraining factor for image convergence.

5.4.3 Polarization

The difference the inclusion of polarization makes for highly specular outdoor scenes is shown in figure 5.20. Inter-reflections between skyscrapers, or glass facades and water bodies in outdoor settings can be considerably mispredicted by a non-polarizing renderer. As our model includes in-scattered radiance for finite viewing directions, it also would be capable of reproducing the effect that a part of the in-scattered light in a daytime outdoor scene can be removed with a polarization filter: a part of the haze can be ”lifted” that way.

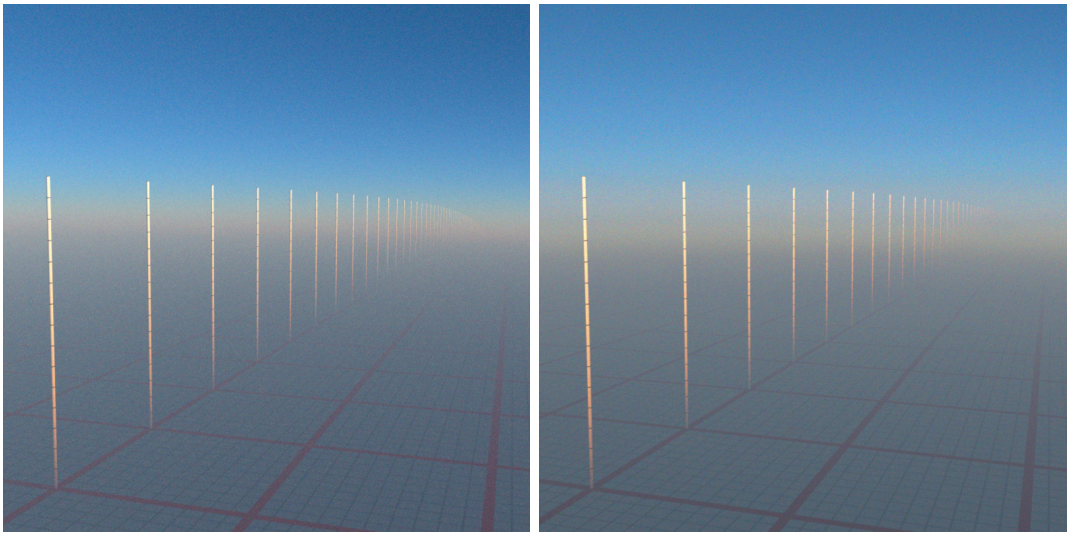

Figure 5.19: **Left:** a brute force rendering of a scene that involves atmospheric perspective. **Right:** The same scene, rendered using the analytical skydome and aerial perspective model. The red squares are 10x10km, and the poles are 15 km high (with stripes every km of altitude). The camera is at 14km altitude, and the sun is at 1 degree above the horizon. Turbidity is 4.

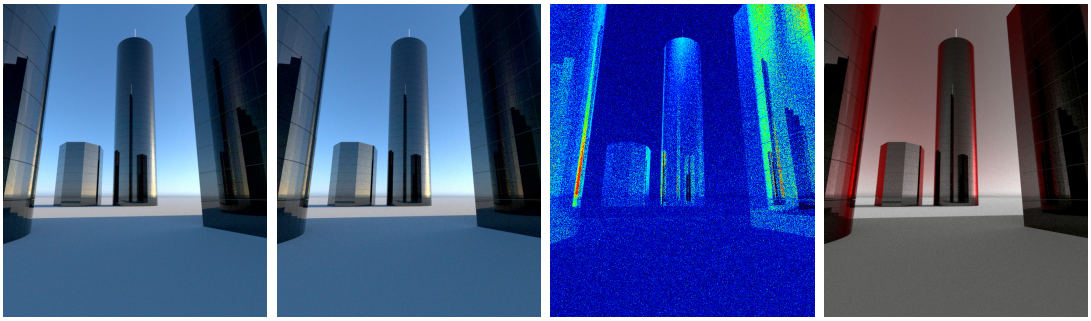

Figure 5.20: A re-creation of the specular architecture scene from Wilkie et al. [2004]: qualitatively similar behavior as seen in that publication can be observed. Image rendered with a polarizing skylight (left), and a non-polarizing version of our model (second from left). The difference image (second from right) shows that skylight polarization not only affects the reflection of the sky in the building facades, but also the inter-reflections of the buildings. The image on the right shows the degree of polarization for 550 nm.

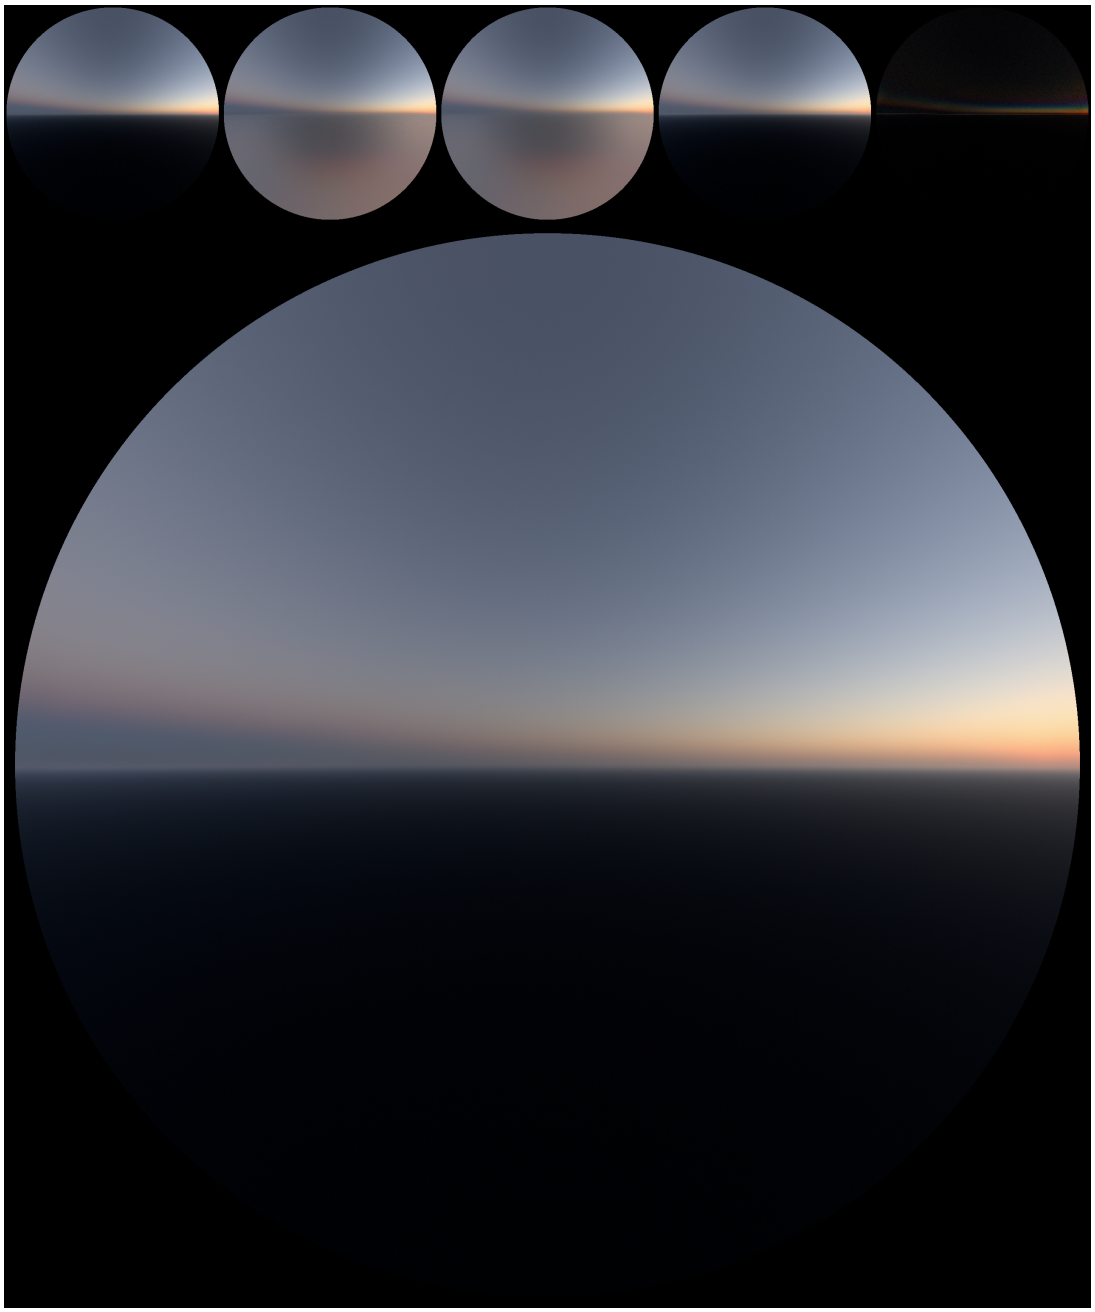

Figure 5.21: The fitting process. Solar elevation $-2.43°$, altitude 405 m. Top row, left to right: the input image, pre-emphasized input image, tensor decomposition output, final de-emphasized output, error. Bottom image is a large rendering of the final output.

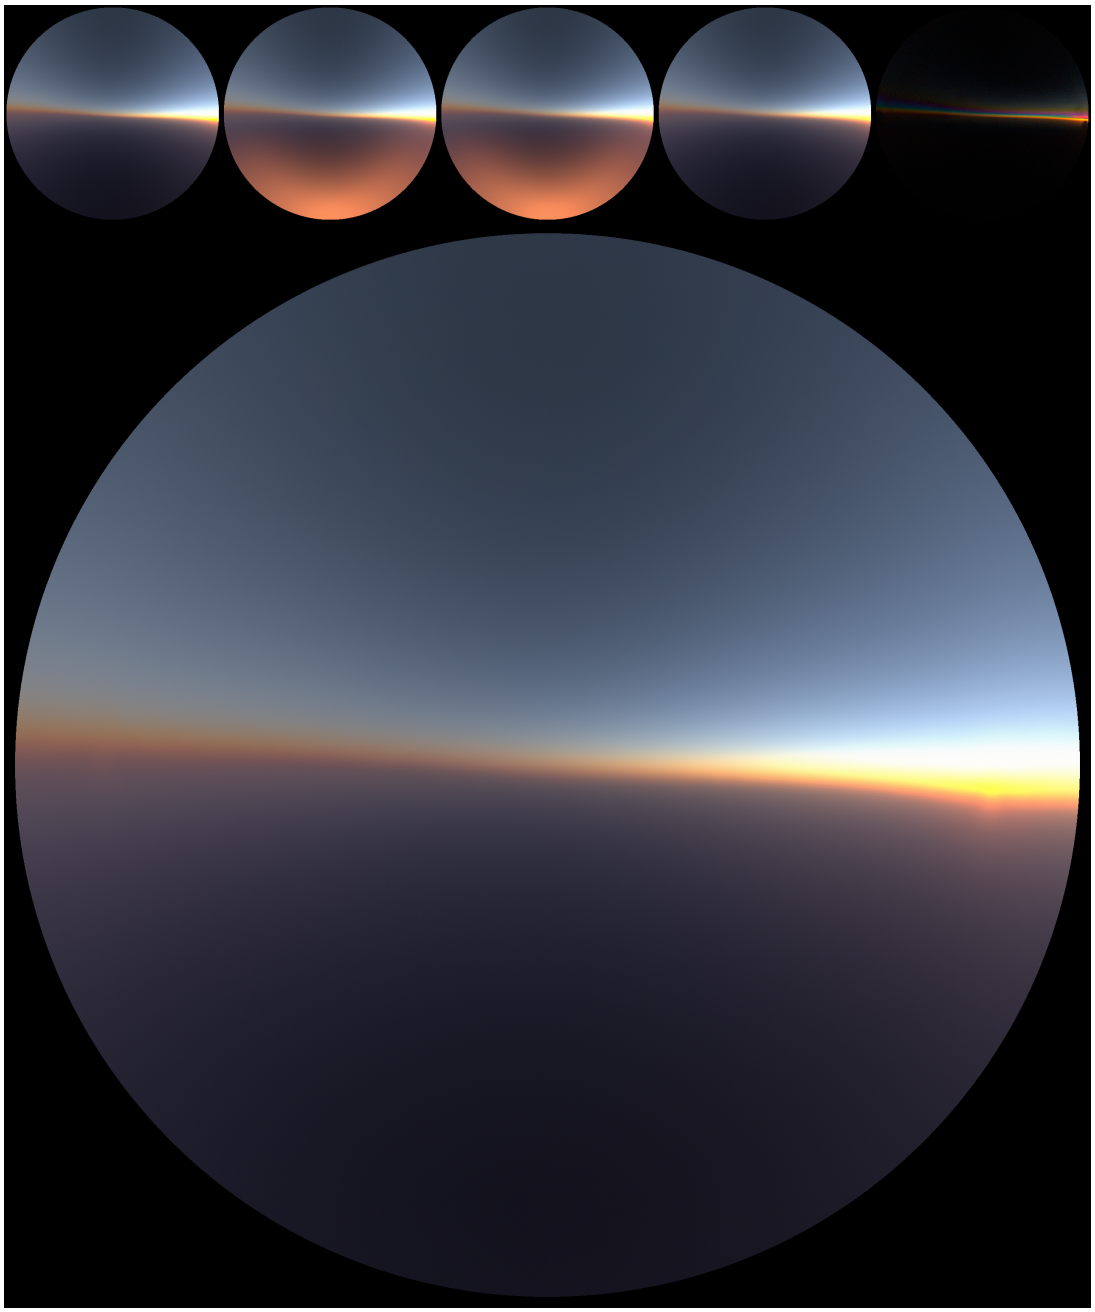

Figure 5.22: The fitting process. Solar elevation $-2.43°$, altitude 15 km. Top row, left to right: the input image, pre-emphasized input image, tensor decomposition output, final de-emphasized output, error. Bottom image is a large rendering of the final output in a super-circular 220° fisheye projection.

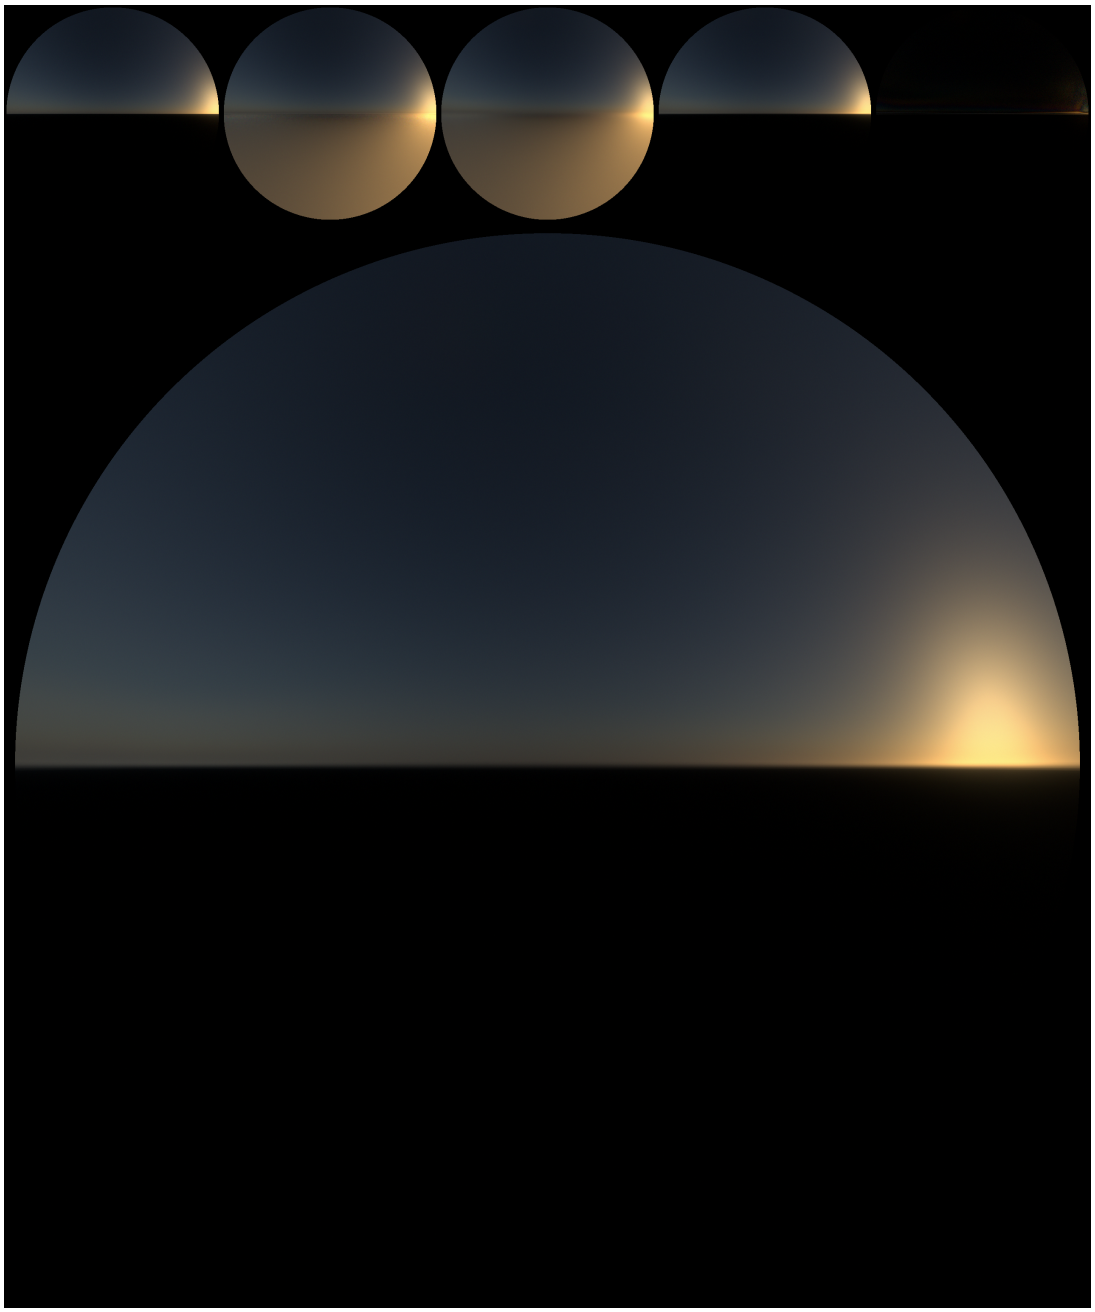

Figure 5.23: The fitting result. Solar elevation 8.2°, altitude 15 m, turbidity 4. The top row shows the intermediate fitting stages. The large image is the final model output. The image is intentionally tone-mapped to be dark to show the detail of the aureole.

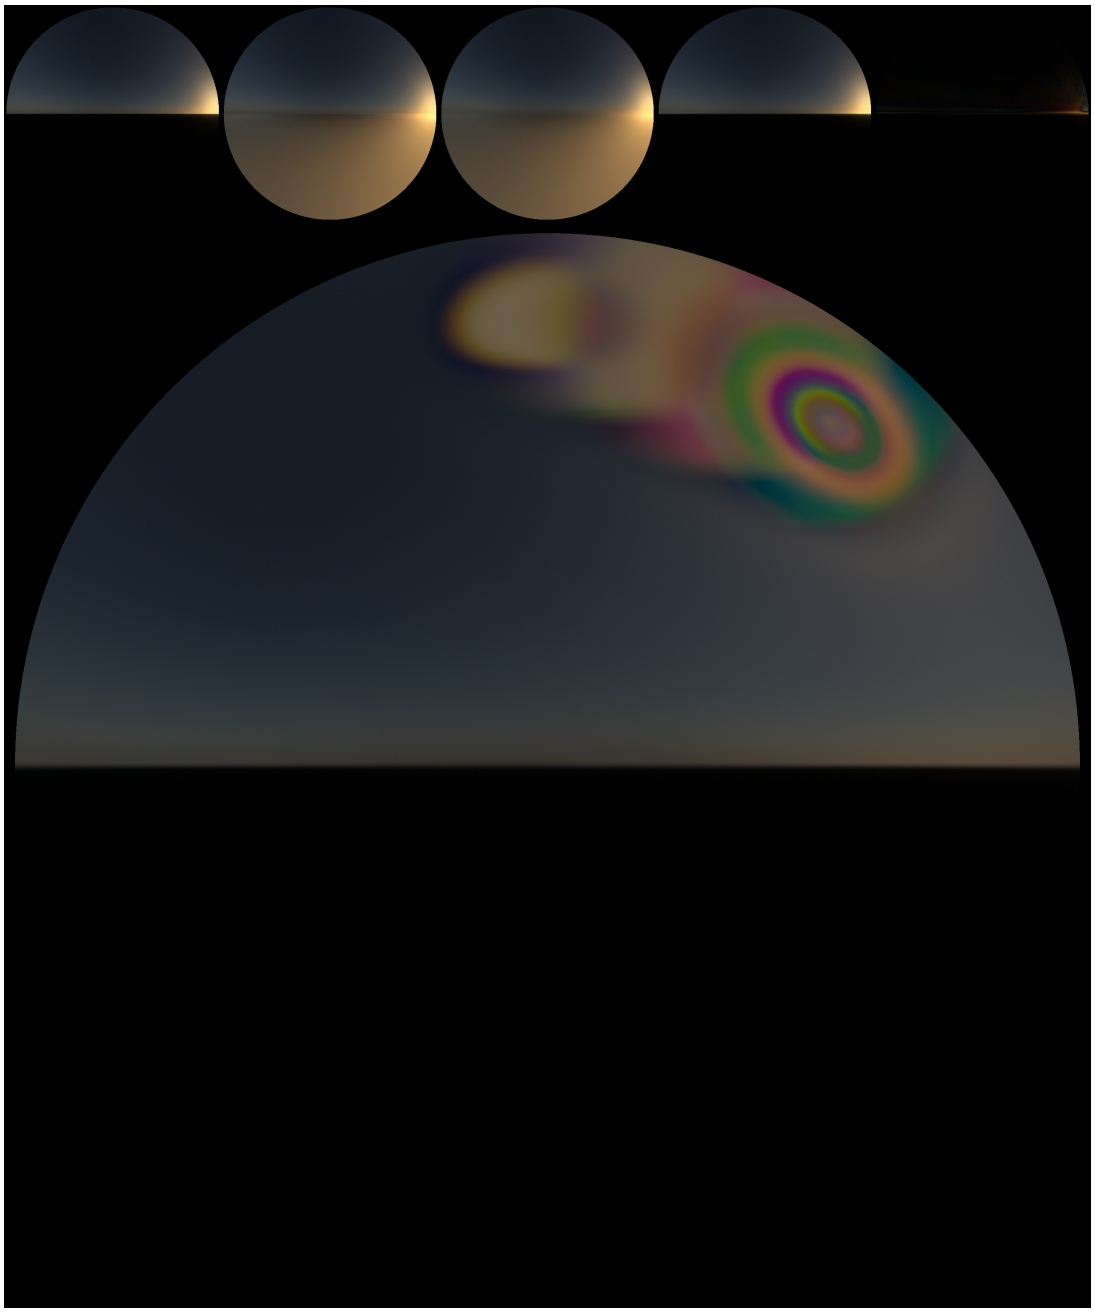

Figure 5.24: A reprojection experiment. The same fitting data that was used to draw fig. 5.23 is used to draw a skydome at solar elevation 45°, but no tensor in-filling was performed. The image is rendered from an undefined part of the tensor, which produces the rainbow artifact at the spot where the sun should be.

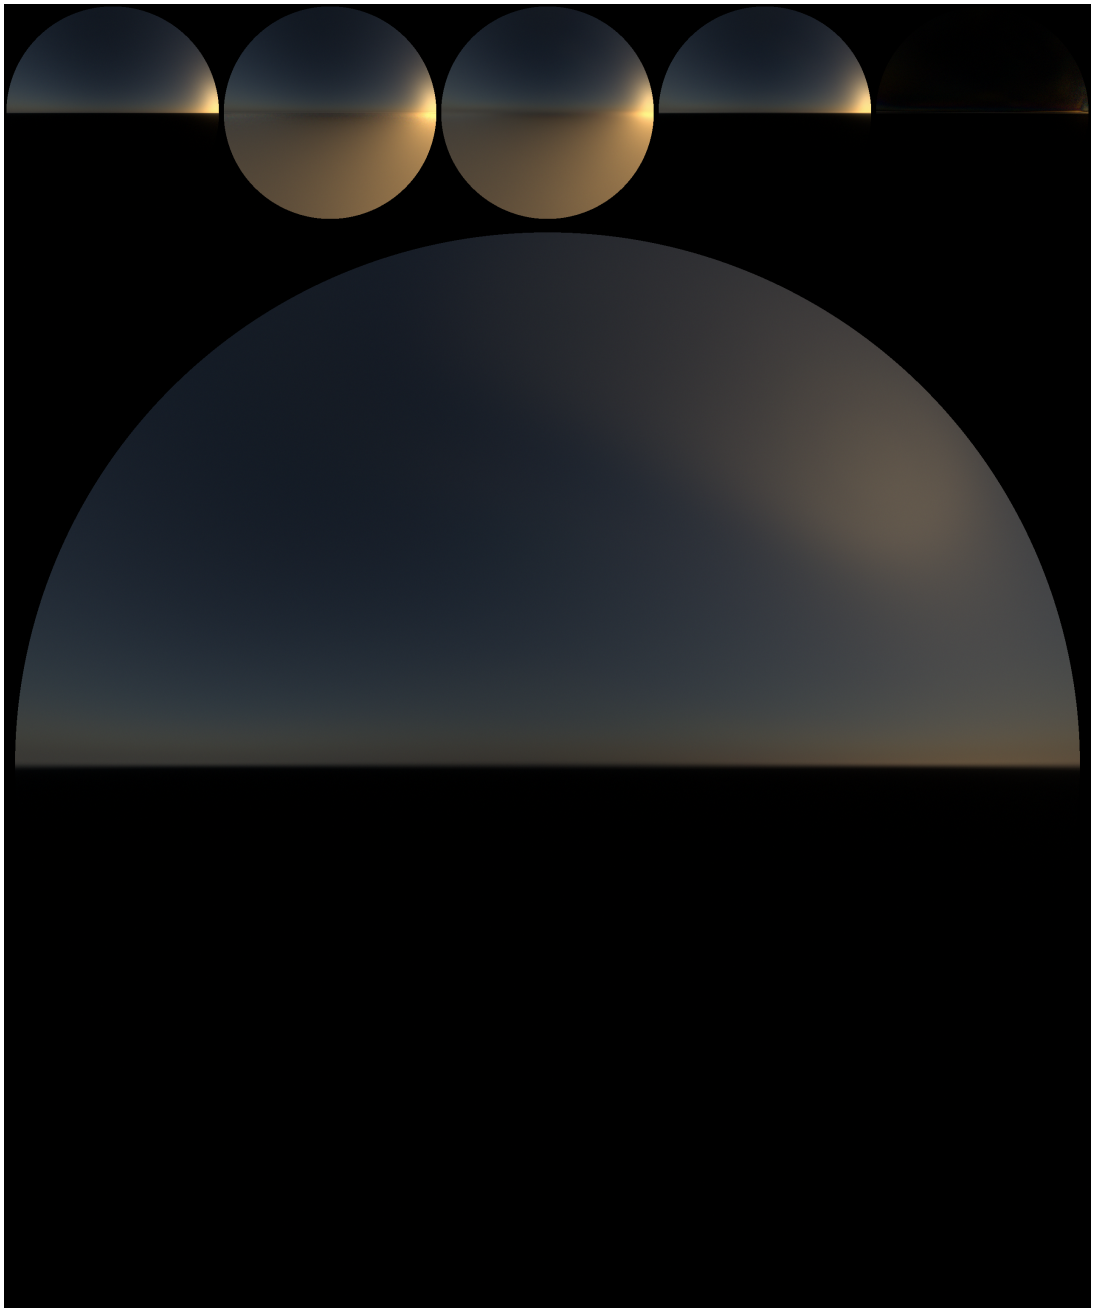

Figure 5.25: The same reprojection experiment as the one showcased on fig. 5.24, but this time with an in-filled tensor. The original fitting was done at solar elevation 8°, the reprojection places the sun at 45°.

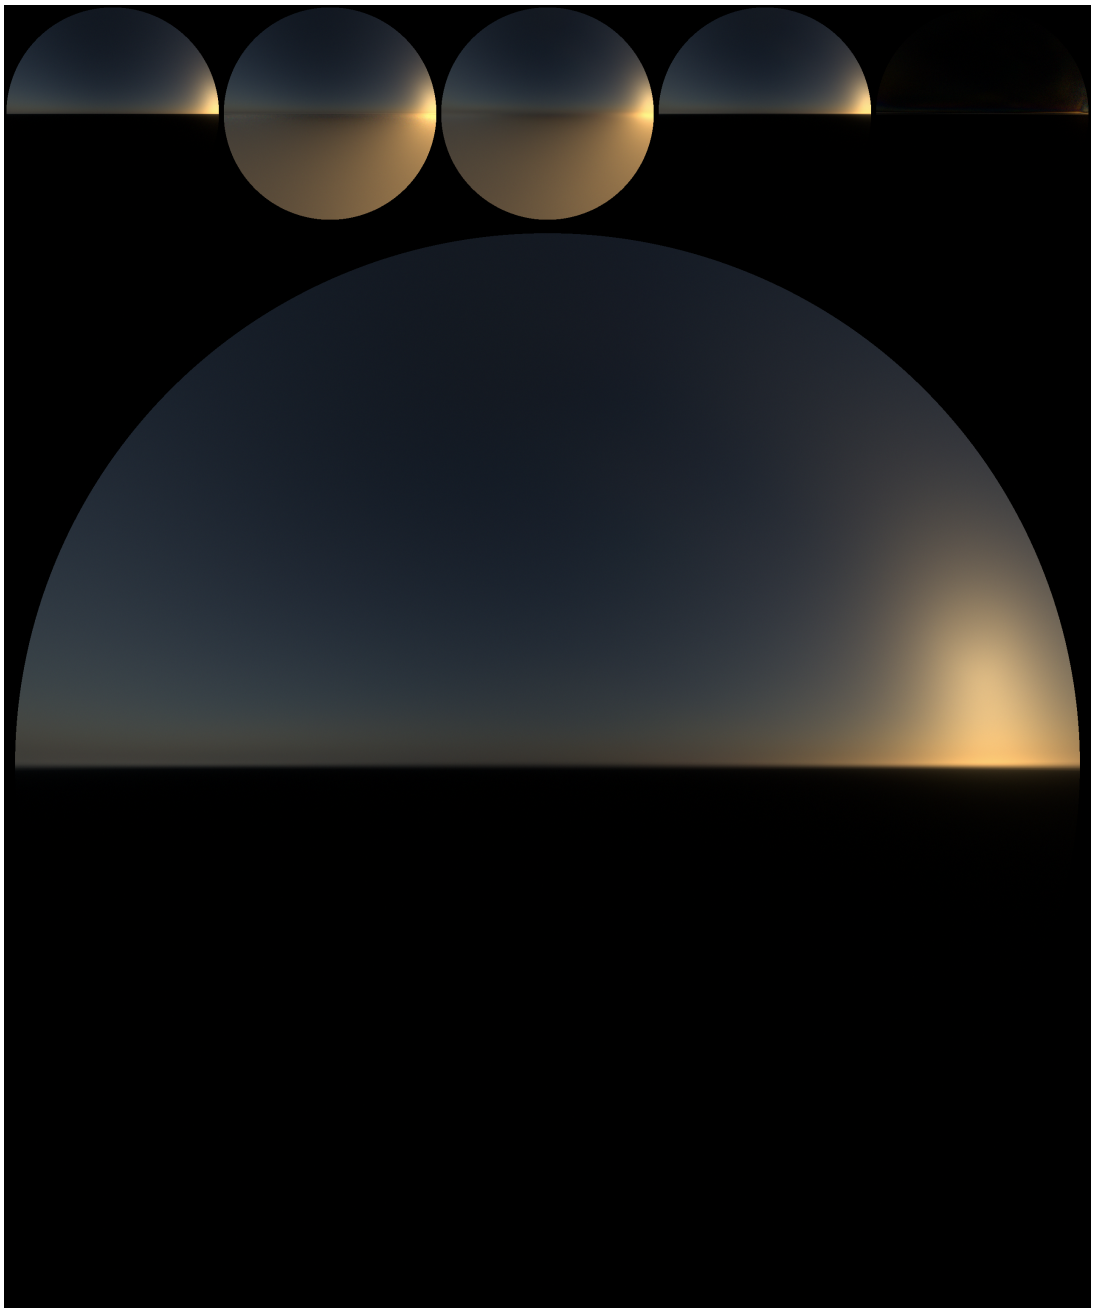

Figure 5.26: A more realistic reprojection. From the original 8° to 16°.

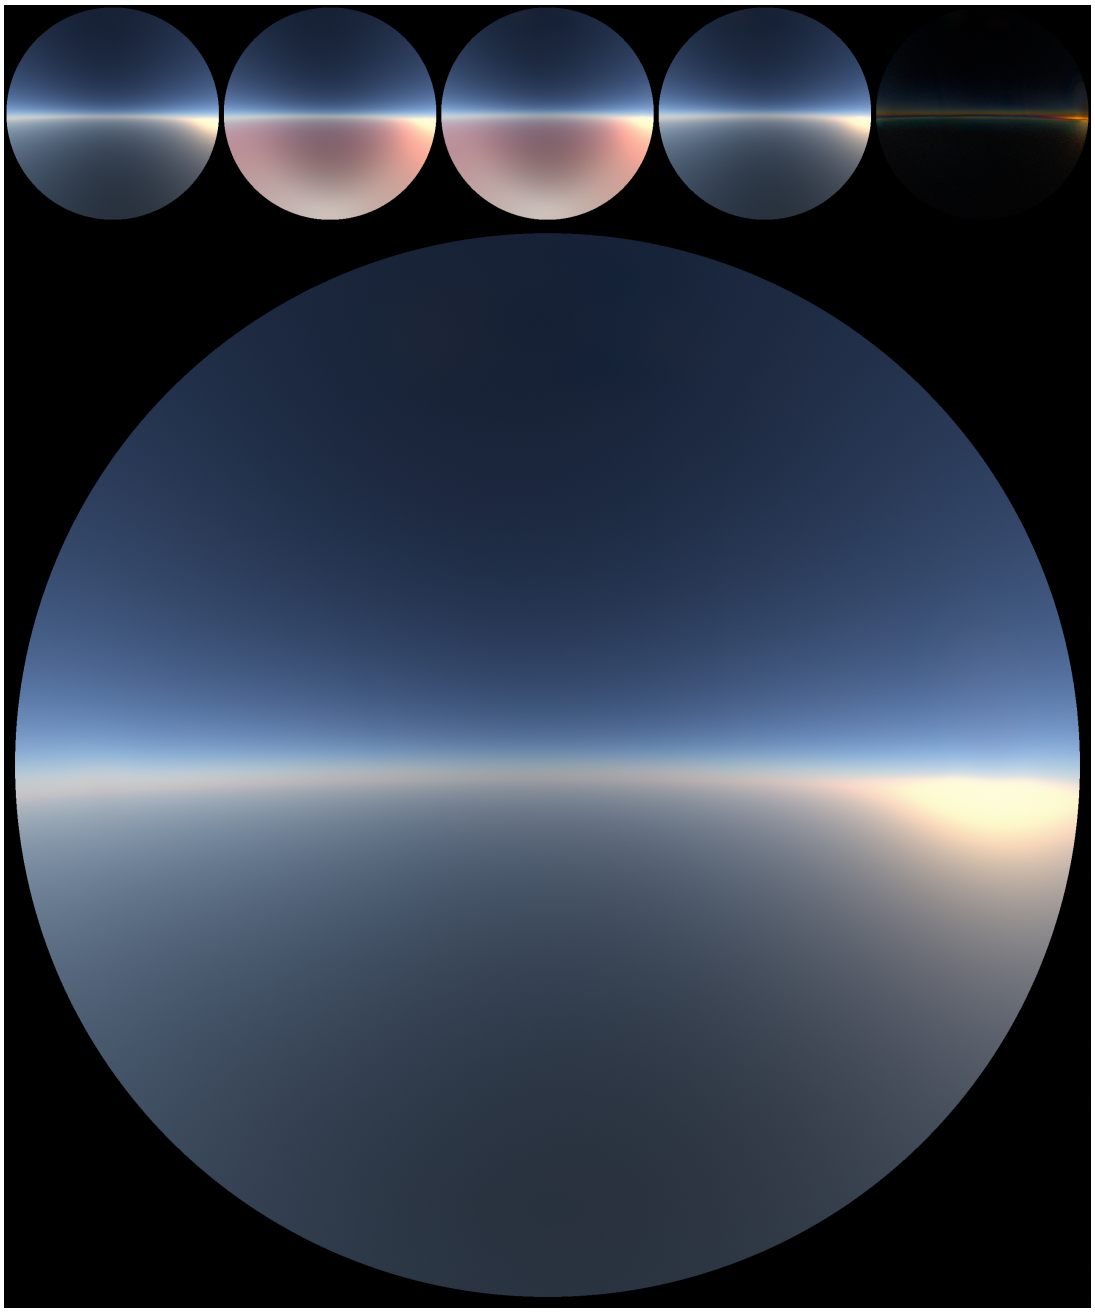

Figure 5.27: The fitting result. Solar elevation 8.2°, altitude 15 km, turbidity 4. The top row shows the intermediate fitting stages. The large image is the final model output. This is the kind of view visible from an airplane at cruising altitude, watching the sunset.

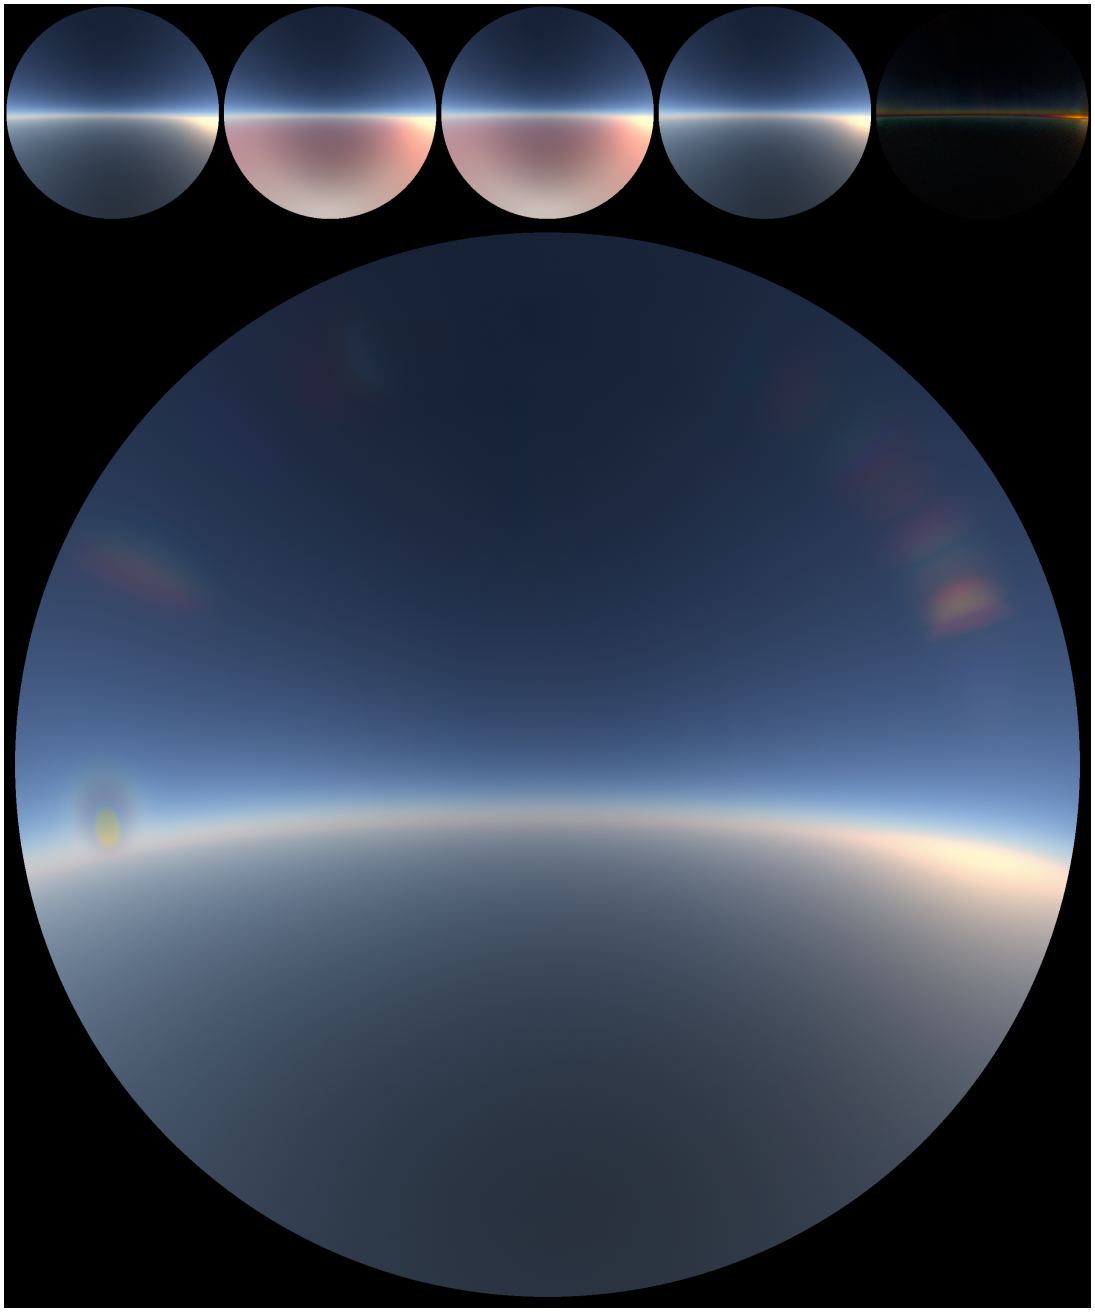

Figure 5.28: A reprojection experiment. The fitting data from fig. 5.27 was used to draw an image at altitude 300 km. The Earth's curvature is apparent from this altitude. Artifacts produced by not in-filling the tensors.

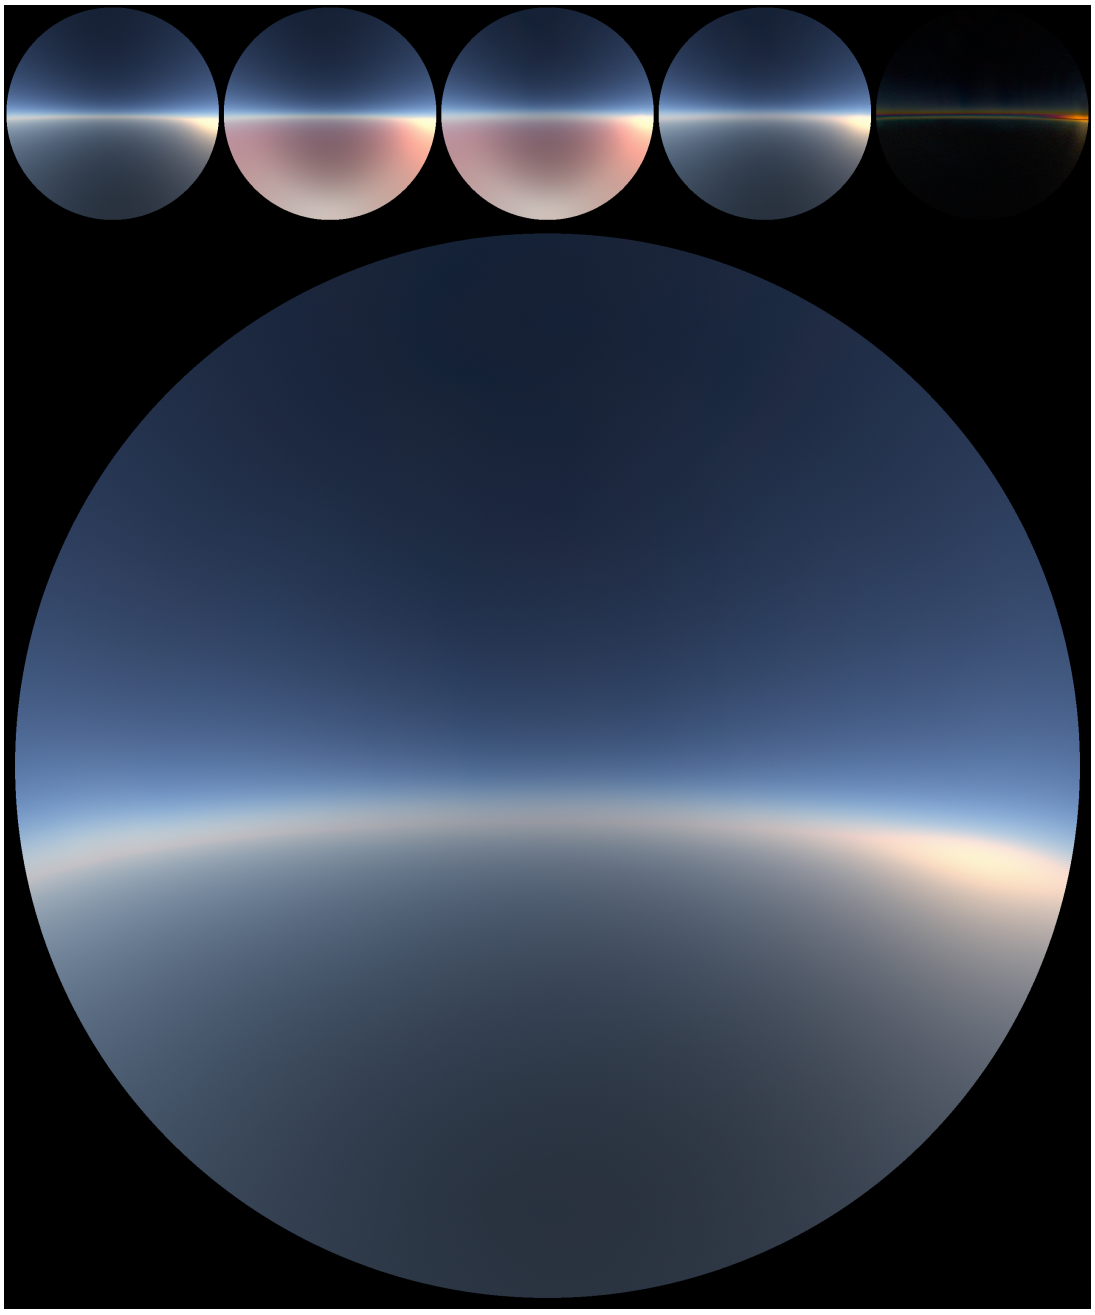

Figure 5.29: The same reprojection experiment attempted in fig. 5.28, but this time with in-filled tensors, which fixes the artifacts.

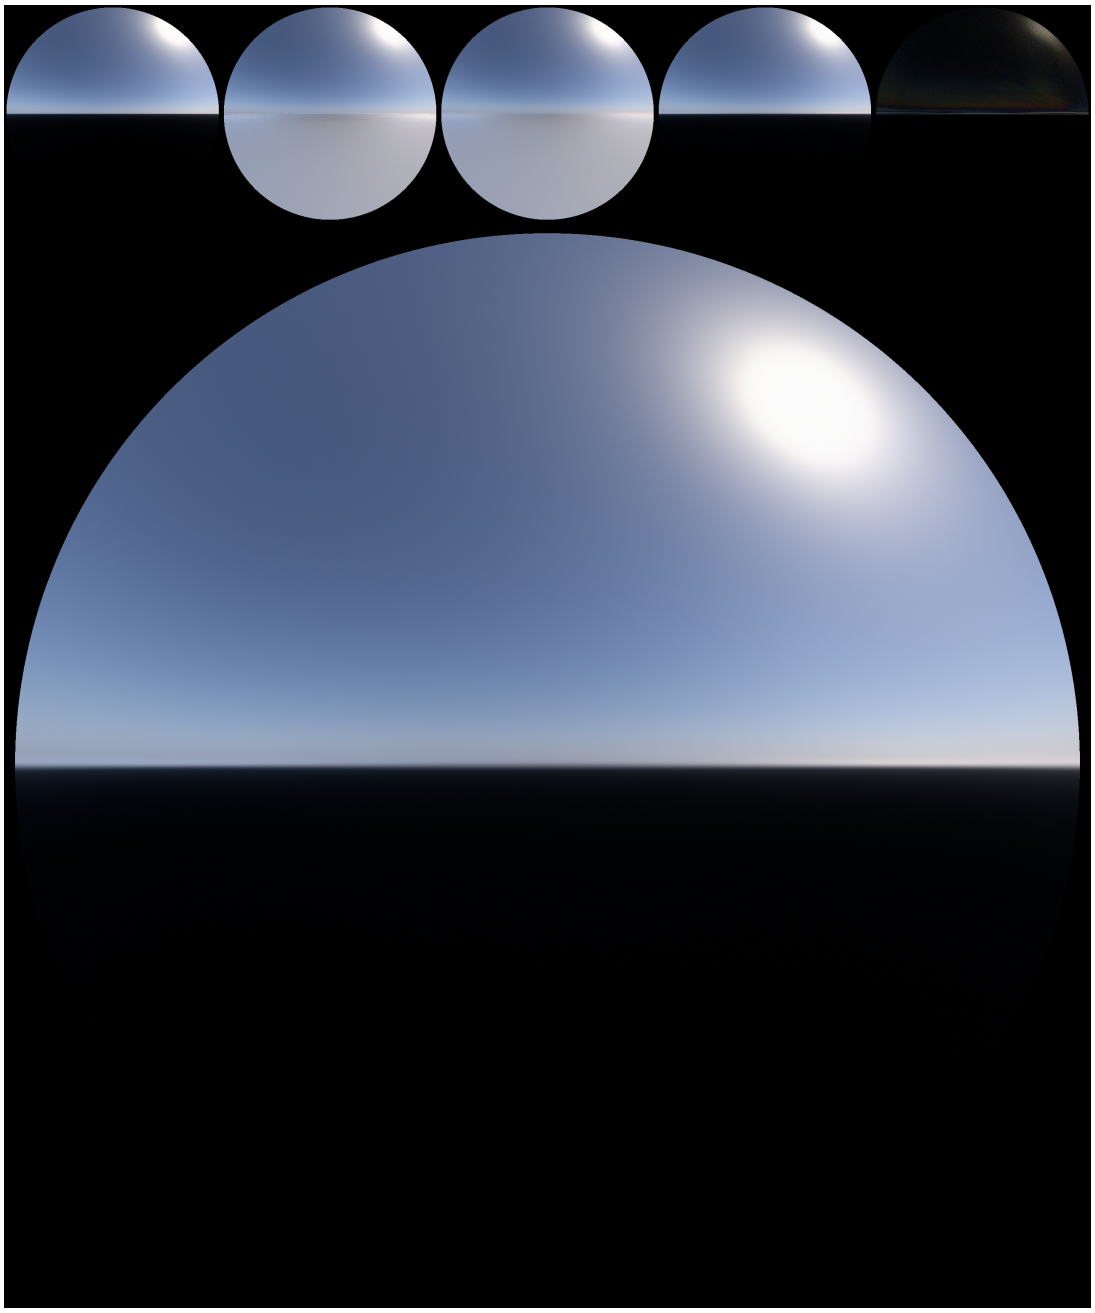

Figure 5.30: The fitting result. Solar elevation 55°, altitude 15 m, turbidity 4.

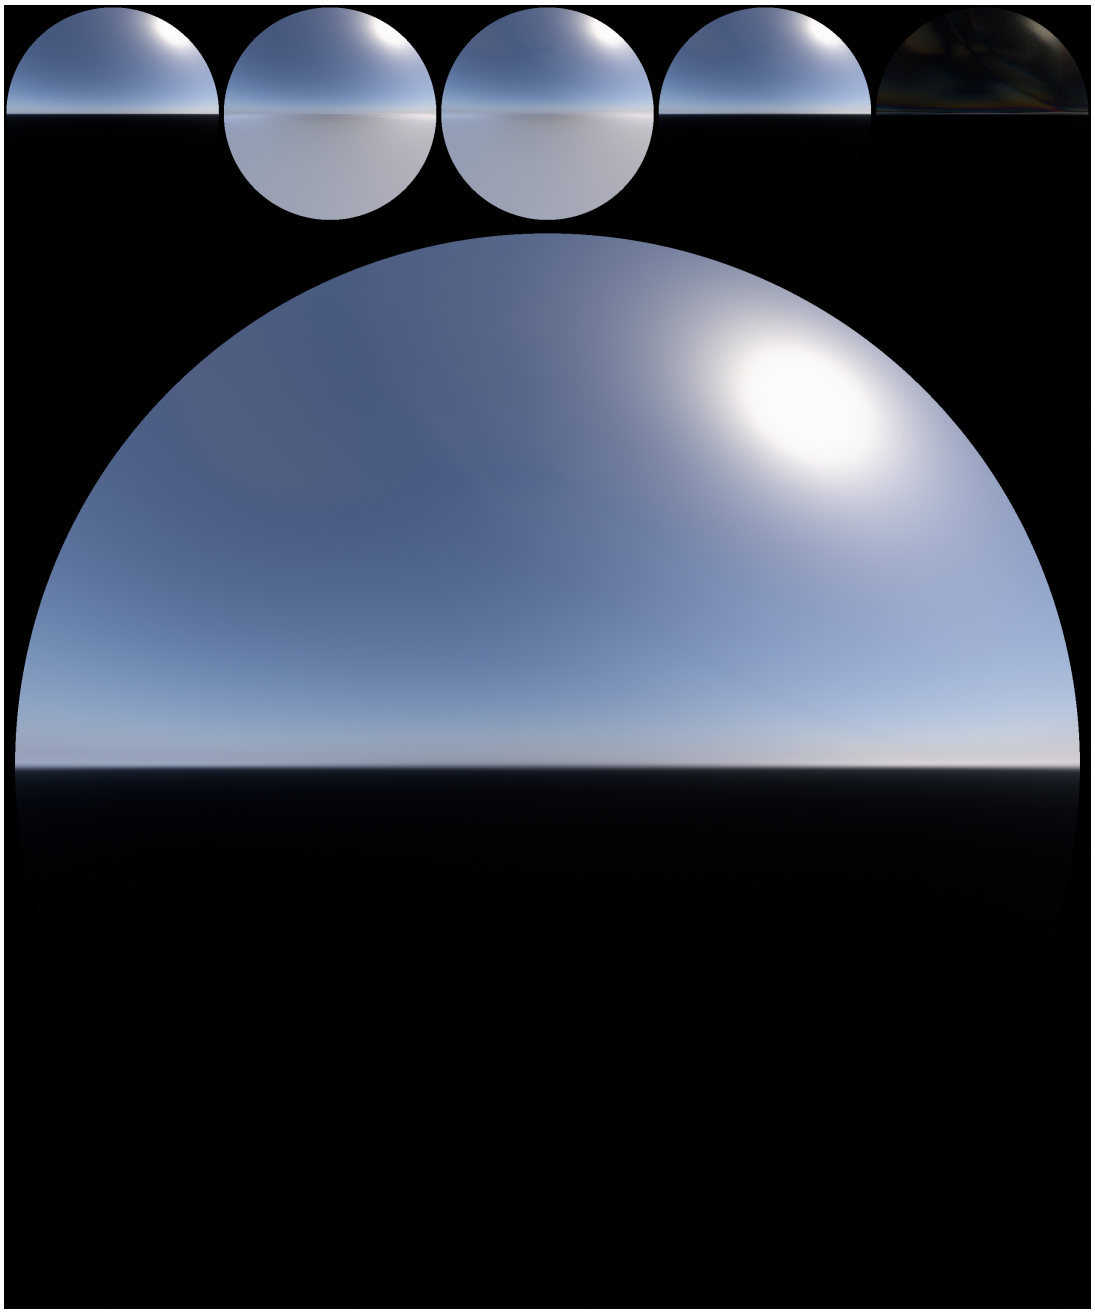

Figure 5.31: Same as fig. 5.30, but without the tensor filtering, which produces an artifact appearing as a ring of whiter color at approx. 45° around the sun.

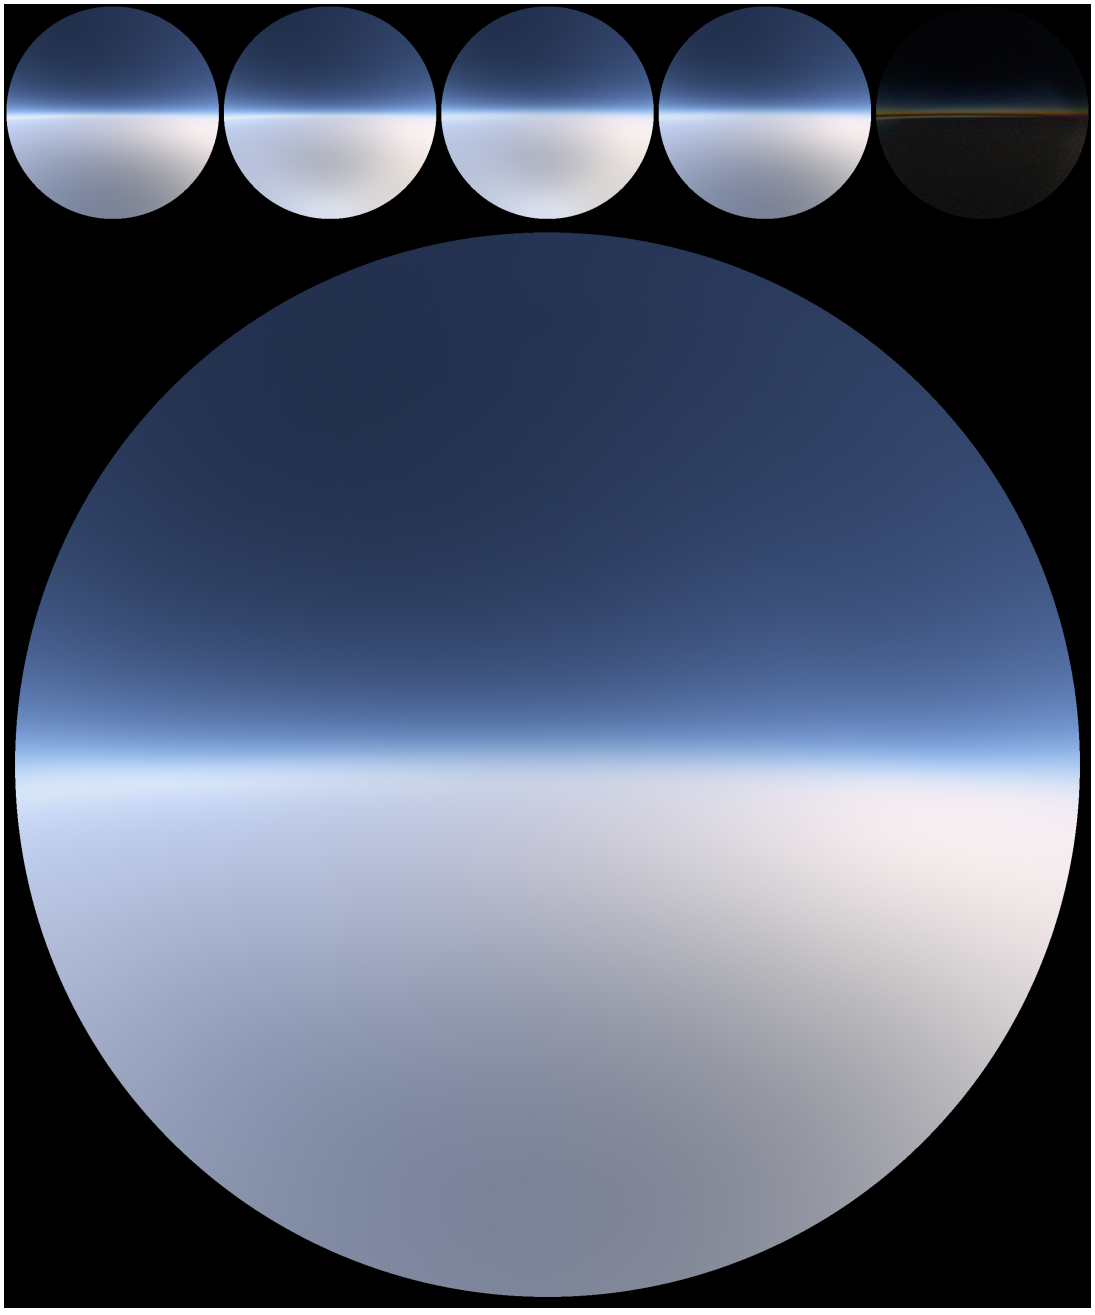

Figure 5.32: The fitting result. Solar elevation 55°, altitude 15 km, turbidity 11. Flying high above a cloud layer. No aureole is visible because the camera is above the aerosol layer – there's no medium in which the aureole would form.

Figure 5.33: A sunset photo and its recreation using the analytic model.

6. Conclusion

In this thesis we have described two skydome models, each one presented as a complete integrated package.

In chapter 4, we have presented a model that can be used as a drop-in replacement for the widely used Preetham model and which provides improved fidelity. The model retains the desirable properties of the Preetham model: it is a very simple function that can be evaluated in real time and is configured with a small and comprehensive set of parameters. The model is fully spectral and data is also provided for the near-UV band, making it useful for predicting the appearance of fluorescent objects (such as paint with optical brighteners). The model has already seen widespread adoption in software packages such as the `Blender` toolset, and the `LuxCoreRender`, `Mitsuba`, `Manuka` and `Corona` renderers.

The second and improved model, described in chapter 5, presents a significant departure from the approach used in our previous model and the models of Perez and Preetham. A novel way of fitting the reference data using tensor decomposition is presented. The resulting analytical model provides a better fidelity while still being easy to implement and integrate and usable in real time. In addition to being fully spectral and providing near-UV data, the new model also provides polarization data, which makes it useful for accurately predicting the appearance of the reflection of the sky on glossy surfaces such as water and glass. The model is based on reference data generated by a path-tracer that is state of the art with regards to physically modeling light interaction with the atmosphere. It is, however, worth noting that although the new model provides better quality outputs, it does not completely obsolete the old one because for some applications the simplicity of the older model might be more important than the additional features of the newer model.

Acknowledgements

This research has been partially funded by GAČR grant 16-08111S.

Bibliography

Eric Bruneton. A qualitative and quantitative evaluation of 8 clear sky models. *CoRR*, abs/1612.04336, 2016. URL `http://arxiv.org/abs/1612.04336`.

Eric Bruneton and Fabrice Neyret. Precomputed atmospheric scattering. In *Proceedings of the Nineteenth Eurographics Conference on Rendering*, EGSR '08, pages 1079–1086, Aire-la-Ville, Switzerland, Switzerland, 2008. Eurographics Association.

W. A. Coleman. Mathematical verification of a certain Monte Carlo sampling technique and applications of the technique to radiation transport problems. *Nuclear Science and Engineering*, 32(1):76–81, 1968.

William M. Cornette and Joe Shanks. Physically reasonable analytic expression for the single-scattering phase function. *Applied optics*, 31:3152–60, 06 1992. doi: 10.1364/AO.31.003152.

Kinsell L. Coulson. *Solar and terrestrial radiation : methods and measurements / by Kinsell L. Coulson.* Academic Press, New York :, 1975. ISBN 0121929507.

Les Cowley and Michael Schoeder. Halosim. `https://www.atoptics.co.uk/halo/halfeat.htm`, 2004. Accessed: 2019-05-28.

Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, Sep 1936. ISSN 1860-0980.

R. G. Eldridge. A few fog drop-size distributions. *Journal of the Atmospheric Sciences*, 18(5):671–676, 1961.

Oskar Elek and Petr Kmoch. Real-time spectral scattering in large-scale natural participating media. In Helwig Hauser and Reinhard Klein, editors, *Proceedings of Spring Conference on Computer Graphics 2010*, pages 83–90. Comenius University, Bratislava, 2010.

Claudia Emde, Robert Buras-Schnell, Arve Kylling, Bernhard Mayer, Josef Gasteiger, Ulrich Hamann, Jonas Kylling, Bettina Richter, Christian Pause, Timothy Dowling, and Luca Bugliaro. The libradtran software package for radiative transfer calculations (version 2.0.1). *Geoscientific Model Development*, 9:1647–1672, 05 2016.

V. Gorshelev, A. Serdyuchenko, M. Weber, W. Chehade, and J. P. Burrows. High spectral resolution ozone absorption cross-sections – Part 1: Measurements, data analysis and comparison with previous measurements around 293 K. *Atmospheric Measurement Techniques*, 7(2):609–624, 2014.

Sven-Erik Gryning and Ekaterina Batchvarova. *Air Pollution Modeling and Its Application XIII.* Springer Science & Business Media, 2012.

Ralf Habel, Bogdan Mustata, and Michael Wimmer. Efficient spherical harmonics lighting with the preetham skylight model. In Katerina Mania and Erik Reinhard, editors, *Eurographics 2008 - Short Papers*, pages 119–122. Eurographics Association, Apr 2008.

Jörg Haber, Marcus A. Magnor, and Hans-Peter Seidel. Physically-based simulation of twilight phenomena. *ACM Trans. Graph.*, 24:1353–1373, 2005.

L. G. Henyey and J. L. Greenstein. Diffuse radiation in the Galaxy. *Astrophys. J.*, 93:70–83, 1941.

Michael Hess, Peter Koepke, and I Schult. Optical properties of aerosols and clouds: The software package opac. *Bulletin of the American Meteorological Society*, 79:831–844, 05 1998.

Lukáš Hošek and Alexander Wilkie. An analytic model for full spectral sky-dome radiance. *ACM Trans. Graph.*, 31(4):95:1–95:9, July 2012. ISSN 0730-0301.

Lukáš Hošek and Alexander Wilkie. Adding a solar-radiance function to the hošek-wilkie skylight model. *IEEE Comput. Graph. Appl.*, 33(3):44–52, May 2013. ISSN 0272-1716.

E. O. Hulburt. Explanation of the brightness and color of the sky, particularly the twilight sky. *J. Opt. Soc. Am.*, 43:113–118, 1953.

ISO 2533:1975. Standard atmosphere. Standard, International Organization for Standardization, Geneva, CH, March 1975.

Wojciech Jarosz, Derek Nowrouzezahrai, Robert Thomas, Peter-Pike Sloan, and Matthias Zwicker. Progressive photon beams. *ACM Trans. Graph. (SIGGRAPH Asia 2011)*, 30(6), 2011.

AK Kaifel, M Felder, C DeClercq, and J-C Lambert. New dynamic nnorsy ozone profile climatology. *Atmospheric Measurement Techniques Discussions*, 5(1): 775–812, 2012.

James T. Kajiya. The rendering equation. *SIGGRAPH Comput. Graph.*, 20(4): 143–150, August 1986. ISSN 0097-8930.

Joseph T. Kider, Jr., Daniel Knowlton, Jeremy Newlin, Yining Karl Li, and Donald P. Greenberg. A framework for the experimental comparison of solar and skydome illumination. *ACM Trans. Graph.*, 33(6):180:1–180:12, November 2014. ISSN 0730-0301.

T. Kolda and B. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009.

Peter Kutz. Sky Renderer project blog. `http://skyrenderer.blogspot.com`, 2013. Accessed: 2015-12-31.

Eric P. Lafortune and Yves D. Willems. Rendering participating media with bidirectional path tracing. In Xavier Pueyo and Peter Schröder, editors, *Rendering Techniques '96*, pages 91–100, Vienna, 1996. Springer Vienna. ISBN 978-3-7091-7484-5.

Raymond L. Lee, Wolfgang Meyer, and Goetz Hoeppe. Atmospheric ozone and colors of the Antarctic twilight sky. *Applied Optics*, 50(28):162–171, 2011.

B Maheu, G Gouesbet, and G Grehan. A concise presentation of the generalized lorenz-mie theory for arbitrary location of the scatterer in an arbitrary incident profile. *Journal of Optics*, 19(2):59–67, mar 1988.

Kresimir Matkovic and Laszló Neumann. Interactive calibration of the mapping of global illumination values to display devices. In *12th Spring Conference on Computer Graphics*, pages 155–161. Comenius University, Bratislava, Slovakia, June 1996. ISBN 80-223-1032-8.

E.J. McCartney. *Optics of the Atmosphere*. Wiley, 1976.

Chris Mclinden. Observations of atmospheric composition from nasa er2 spectroradiometer measurements. 01 1998.

Victor Meyer-Rochow and Rüdiger Wehner. Polarization patterns of the summer sky and its neutral points measured by full-sky imaging polarimetry in finnish lapland north of the arctic circle. *Proceedings of The Royal Society A: Mathematical, Physical and Engineering Sciences*, 457:1385–1399, 06 2001.

Bailey Miller, Iliyan Georgiev, and Wojciech Jarosz. A null-scattering path integral formulation of light transport. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, 38(4), jul 2019.

Hans Mueller. The foundations of optics. In *Proceedings of the Winter Meeting of the Optical Society of America*. OSA, 1948.

H. Neckel and D. Labs. Solar limb darkening 1986-1990 (lambda lambda 303 to 1099nm). *Solar Physics*, 153:91–114, August 1994. ISSN 0038-0938.

Tomoyuki Nishita, Takao Sirai, Katsumi Tadamura, and Eihachiro Nakamae. Display of the earth taking into account atmospheric scattering. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '93, pages 175–182, New York, NY, USA, 1993. ACM. ISBN 0-89791-601-8.

Mark Pauly, Thomas Kollig, and Alexander Keller. Metropolis light transport for participating media. In Bernard Péroche and Holly Rushmeier, editors, *Rendering Techniques 2000*, pages 11–22, Vienna, 2000. Springer Vienna. ISBN 978-3-7091-6303-0.

R. Perez, R. Seals, and J. Michalsky. All-weather model for sky luminance distribution–preliminary configuration and validation. *Solar Energy*, 50(3): 235 – 245, 1993. ISSN 0038-092X.

A. J. Preetham, Peter Shirley, and Brian Smits. A practical analytic model for daylight. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '99, pages 91–100, New York, NY, USA, 1999. ACM Press/Addison-Wesley Publishing Co. ISBN 0-201-48560-5.

Matthias Raab, Daniel Seibert, and Alexander Keller. Unbiased global illumination with participating media. In Alexander Keller, Stefan Heinrich, and Harald Niederreiter, editors, *Monte Carlo and Quasi-Monte Carlo Methods 2006*, pages 591–605. Springer Berlin Heidelberg, 2008. ISBN 978-3-540-74496-2.

John William Strutt. *Scientific Papers*. Cambidge University Press Warehouse, Cambridge, 1899.

Eric Veach and Leonidas J. Guibas. Optimally combining sampling techniques for monte carlo rendering. In *Proceedings of the 22Nd Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '95, pages 419–428, New York, NY, USA, 1995. ACM. ISBN 0-89791-701-4.

Eric Veach and Leonidas J. Guibas. Metropolis light transport. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '97, pages 65–76, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co. ISBN 0-89791-896-7.

Xin Wang, Jun Gao, Zhiguo Fan, and Nicholas W Roberts. An analytical model for the celestial distribution of polarized light, accounting for polarization singularities, wavelength and atmospheric turbidity. *Journal of Optics*, 18(6): 065601, 2016.

A. Wilkie, S. Nawaz, M. Droske, A. Weidlich, and J. Hanika. Hero wavelength spectral sampling. In *Proceedings of the 25th Eurographics Symposium on Rendering*, EGSR '14, pages 123–131, Aire-la-Ville, Switzerland, Switzerland, 2014. Eurographics Association. doi: 10.1111/cgf.12419. URL `http://dx.doi.org/10.1111/cgf.12419`.

Alexander Wilkie and Andrea Weidlich. Polarised light in computer graphics. In *SIGGRAPH Asia 2012 Courses*, 2012.

Alexander Wilkie, Claudia Ulbricht, Robert F. Tobler, Georg Zotti, and Werner Purgathofer. An analytical model for skylight polarization. In Alexander Keller and Henrik Wann Jensen, editors, *Rendering Techniques*, pages 387–398. Eurographics Association, 2004. ISBN 3-905673-12-6.

Alexander Wilkie et al. The Advanced Rendering Toolkit, 2018. https://cgg.mff.cuni.cz/ART/.

Yuk L. Yung. An introduction to atmospheric radiation. by k. n. liou. academic press. second edition, 2002. pp. xiv+583. isbn 0 12 451451 0. *Quarterly Journal of the Royal Meteorological Society*, 129(590):1741–1741, 2003.

Huijie Zhao, Wujian Xu, Ying Zhang, Xudong Li, Hao Zhang, Jiabin Xuan, and Bo Jia. Polarization patterns under different sky conditions and a navigation method based on the symmetry of the aop map of skylight. *Opt. Express*, 26 (22):28589–28603, Oct 2018.

Georg Zotti, Alexander Wilkie, and Werner Purgathofer. A critical review of the Preetham skylight model. In Vaclav Skala, editor, *WSCG ' 2007 Short Communications Proceedings I*, pages 23–30. University of West Bohemia, 2007. ISBN 978-80-86943-02-2.

List of publications

- Hošek, Lukáš and Wilkie, Alexander, *An Analytic Model for Full Spectral Sky-dome Radiance.* ACM Transactions on Graphics (SIGGRAPH 2012), 31(4), July 2012.

- Hošek, Lukáš and Wilkie, Alexander, *Adding a Solar-Radiance Function to the Hošek-Wilkie Skylight Model.* IEEE Computer Graphics and Applications, 33(4), May 2013.

- Wilkie, Alexander and Hošek, Lukáš, *Predicting Sky Dome Appearance on Earth-like Extrasolar Worlds.* Proceedings of the 29th Spring Conference on Computer Graphics (SCCG), 2013.

A. Attachments

A.1 The Atmosphere Description Format (ADF)

This section provides a description of the Atmosphere Description Format (ADF), which is internally being used by the reference renderer described in chapter 5. Each ADF dataset is a single ASCII encoded file, describes one particular static atmospheric set-up, and consists of a number of *blocks*. Each block contains information about how a particular property of the atmosphere changes as height increases.

There are three types of block, each of which can occur an arbitrary number of times, in arbitrary order. The types are *Rayleigh scatterer*, *Mie scatterer*, and *absorber*.

In addition to data specific for each type of block, each of these always contains a tabulated height profile which expresses how the molar concentration of the block contents changes with altitude. The format is intentionally a bare-bones format which just contains the actual numerical parameters that are directly used by the Monte Carlo simulations in `atmo_sim` and `haze_sim` (a modified version of `atmo_sim` that generates reference transmittance data). ADF files are normally generated by `adf_gen`, a tool that implements the synthetic profile approach described in section 5.1.1, according to user specifications: for a given turbidity value, and a given technique to determine Mie scattering properties, `adf_gen` will output an ADF file that has the desired height profiles. However, as the files are human-readable ASCII, they could conceivably also be filled out by hand.

A.1.1 File Header

The first line of an ADF file is always a static header line, followed by a comment which stores the turbidity value for which `adf_gen` generated this particular ADF file. This is in turn followed by a curly brace block which contains information about the planet: the radius of the planet [m] is given, as well as the apparent diameter [rad] of the solar disc.

```
#Atmosphere Description Format V1.1
#turbidity 4.0
{
        planet radius = 6378000.0
        solar diameter = 0.008727
}
```

A.1.2 ADF Blocks

A description of each type of block that can be found in an ADF file follows. Generally, the first line of a block consists of a hash sign, followed by a line with identifying information about to the particular block, for example `#Ozone`. This line only serves for information purposes, and is ignored during the actual simulation.

This is then followed by a block that is surrounded by curly brackets. The next line contains the name of the interaction type. This is followed by a line of the form `Step = X`, which specifies vertical step size of the quantity in meters. For example, `Step = 1000` would tabulate the property in question every 1000 meters.

This is then followed by varying properties (see below) that are specific for the type of interaction (Rayleigh, Mie, absorber). Finally, the tabulated values of molar concentration [$\mathrm{mol} \cdot \mathrm{m}^{-3}$] at increasing altitudes are listed, separated by commas, and surrounded by curly brackets. Additional properties, such as the "extent" of the atmosphere can be inferred from the file format, i.e. the largest values of step $\times$ number of table values, over all blocks.

It is also worth noting that there is a reason for the curly brace text format found in ADF files: it makes it easy to copy & paste file contents into Mathematica notebooks for plotting purposes.

Rayleigh Scatterer

For Rayleigh scattering, the block contains the depolarization factor `depolarization` and the volume polarizability `polarizability`, followed by the vertical profile of the molar concentration of particles which can undergo Rayleigh scattering in steps of `step` meters.

```
#Rayleigh Scatterer
{
        rayleigh
        step = 10.000000
        depolarization = 0.0279
        polarizability = 1.7e-30
        {4.229000 e +01 , 4.225335 e +01 , <... > , 7.259639 e -03}
}
```

Mie Scatterer

For a Mie scatterer, the start of the block also contains information about the layout of spectral parameters typical for Mie. The first parameter, `spectraltype`, is currently always set to regular. This parameter was intended to also allow spectra which are sampled at irregular intervals, but this turned out to be unnecessary so far, and has not been implemented yet.

The next two pairs of values, `lambdastart_ani` and `lambdastep_ani` resp. `lambdastart_cs` and `lambdastep_cs`, describe starting wavelength and sample spacing of the anisotropy parameter and molecular cross-section, respectively. This is followed by two curly brace enclosed lists of these two parameters, and the block is again finished by the vertical concentration profile of the scatterer in question.

The block describes only the scattering property of the particle. Its absorbing property has to be defined separately in an absorber block.

```
#Mie Scatterer (size 1.300000-2.700000 um by 0.025000 um)
{
        mie
        step = 10.000000
        spectraltype = regular
        lambdastart_ani = 320.000000
        lambdastep_ani = 1.000000
        lambdastart_cs = 320
        lambdastep_cs = 40
        10 {0.794174 , 0.794174 , <... > , 0.794174}
        {1.8165 e -12 , 2.2033 e -12 , <... > , 2.4349 e -12}
        {1.739892 e -17 , 1.729871 e -17 , <... > , 1.436144 e -42}
}
```

Absorber

For an absorber such as ozone, or for describing the absorbing property of a Mie scatterer, the block contains wavelength-resolved absorption coefficient data: `lambdastart` and `lambdastep` give start and step values, followed by the data in a curly brace enclosed list. This is followed by vertical profile.

```
#Absorber
{
        absorb
        step = 2500.0
        spectraltype = regular
        lambdastart = 215
        lambdastep = 3
        {9.48 e -23 , 1.38 e -22 , <... > , 6.45 e -27}
        {6.64216 e -7 , 8.3027 e -7 , <... > , 0.0}
}
```

A.2 Atmospheric Transmittance

The following model was developed by Thomas Bashford-Rogers to provide a transmittance function matching the skydome model described in chapter 5.

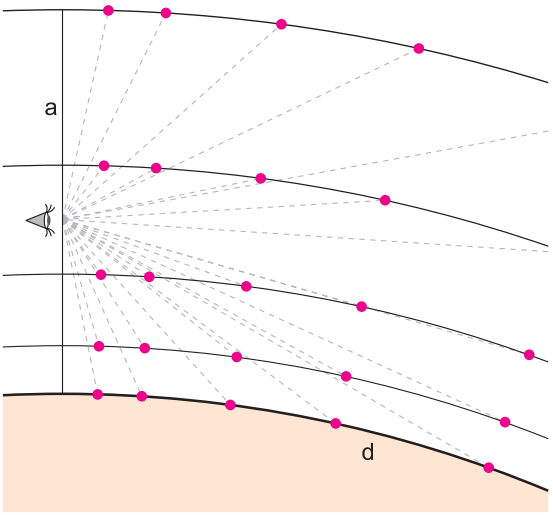

Figure A.1: Illustration of the 2D atmospheric parameterization showing the non-uniformly distributed set of points (pink dots) where transmittance calculated for the model.

The model requires the computation of atmospheric transmission of light between two points for both in-scattered radiance and direct solar radiance. As we assume a rotationally symmetric atmospheric configuration, this reduces to calculating transmission in a 2D atmosphere. The transmission therefore is a 2D distribution which varies with turbidity, wavelength, and altitude of the viewing point. We parameterize the 2D atmosphere by altitude from sea level a and distance along the planet surface d, see Figure A.1. Transmittance is computed using ratio tracking at a non-linearly distributed set of points. These points are non-linearly distributed in the region of the atmosphere which covers the maximum viewable distances from the camera. These points are distributed to capture aspects of the atmosphere such as inversion layers and the transmittance falloff with distance; however any other set of points can be used.

To create a model, we compress the dataset using Singular Value Decomposition (SVD) to produce a low rank approximation to the dataset. We first normalize the rows corresponding to the maximum transmittance at each computed a, leading to a vector of scaling values N. This allows SVD to work in the [0..1] range across that dataset, leading to a reduced rank than would otherwise be required for reconstruction. Secondly, we take the square root of the data to further reduce the effect of inaccuracies in the low rank approximation. Finally, we unroll each normalized and cube rooted 2D slice of the transmittance into a 1D vector, and stack all images corresponding to all configurations (turbidities, wavelengths, and altitudes) into a matrix. Finally, SVD produces a sorted list of eigenvectors and values $\mathbf{U\Sigma V}^*$. We reconstruct from the first R bases $\mathbf{U_R}$, with associated coefficients $\mathbf{C_R} = \mathbf{\Sigma_R V_R^*}$. We store $\mathbf{U_R}$ which is R times the number of non-linearly distributed points for a single atmosphere configuration, $\mathbf{C_R}$ which is R times the number of configurations, and N which is the number

of configurations times the number of points used to store a. In this work, we use $R = 32$, leading to a mean absolute error over the dataset of 0.0022.

At runtime, the transmittance between two points in the atmosphere x_1 and x_2, $\tau(x_1, x_2)$ at a given turbidity and wavelength is computed via reconstructing the transmittance from the reduced rank approximation. Given the altitude, turbidity and wavelength of x_1, the basis coefficients are interpolated $\mathbf{C}'_{\mathbf{R}}$ from the nearest configurations stored in $\mathbf{C}_{\mathbf{R}}$. Then x_2 is projected into the 2D parametrization, leading to a_{x_2} and d_{x_2}, and the inner product between the basis evaluated at a_{x_2} and d_{x_2} and the coefficients is computed. This is then squared, clamped between 0 and 1, and finally correctly scaled by interpolating the scaling values for the configuration at a_{x_2}:

$$\tau(x_1, x_2) = N'(a_{x_2})|\left\langle \mathbf{U_R}(\mathbf{a_{x_2}}, \mathbf{d_{x_2}})|\mathbf{C}'_{\mathbf{R}}\right\rangle^2|_{0,1} \tag{A.1}$$

where $N'(a_{x_2})$ are the interpolated scaling values, and $|\cdot|_{0,1}$ is an operator which clamps a value between 0 and 1.

The model requires the computation of atmospheric transmission of light between two points for both in-scattered radiance and direct solar radiance. As we assume a rotationally symmetric atmospheric configuration, this reduces to calculating transmission in a 2D atmosphere. The transmission therefore is a 2D distribution which varies with turbidity, wavelength, and altitude of the viewing point (a in Figure A.1). We choose to parameterize the 2D distribution via distance D and angle θ for each of the aforementioned atmospheric configurations. The transmittance is computed at a set of quantized distances $D_n = \frac{n}{N}{}^5 S$, where N is the desired quantization level, $n \in [1..N]$, and S is the maximum distance at which transmittance is computed. S was set to $1{,}000km$. We also compute an additional distance set outside the atmosphere, which is required for solar radiance. The entire dataset consists of a 5D table containing 11 turbidities, 11 wavelengths, 21 altitudes, $30 + 1$ quantised distances, at an angular resolution of half a degree, leading to a total of 28.4 million values. Stored at floating point precision, this requires 108 MB memory. S was set to 1000000.

At runtime, the transmission between two points is computed by calculating the altitude of one point, and then the angle and distance d to the second. The transmission is computed by interpolating the values stored in the 5D table. However, linear interpolation of values between the quantized distances can lead to incorrect values, especially in the presence of an inversion layer, as the transmission values can vary significantly between two quantized distances leading to an overestimation of transmittance. Our solution is to interpolate the transmittance values at angles corresponding to the closest altitudes to the point x on the two nearest quantized distances D_n and D_{n+1}. Figure A.1 illustrates this as an orange dashed curve at the altitude of x intersecting the blue curves representing the quantized distances at points p and q. The weights for interpolation are based on Euclidean distances between x, and p and q, for example, the weight corresponding to D_n is $\frac{||q-x||_2}{||q-x||_2+||p-x||_2}$. Figure A.2a shows a reference rendering of transmittance computed via Ratio Tracking, and Figure A.2b is computed using the lookup table.

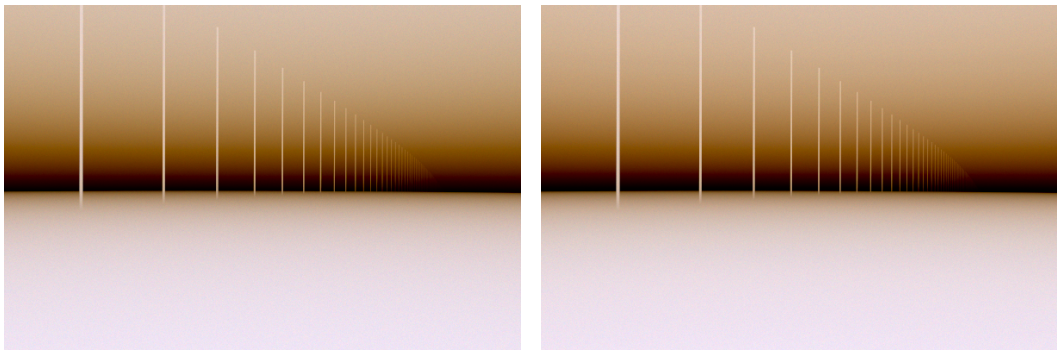

(a) Reference transmittance render (b) SVD Approximation.

Figure A.2: Renderings at an altitude of 8 km. The left image is a reference, and the right is from the SVD approximation.

A.2.1 Finite In-scattered Radiance Calculation

Given that the model provides the ability to calculate radiance for infinite paths and transmittance for finite paths, in-scattered radiance between observer a and viewed point b can be calculated as

$$L_{in}(a,b) = L_{inf}(a,b-a) - \tau(a,b) * L_{inf}(b,b-a) \tag{A.2}$$

where $L_{in}(observer, point)$ denotes in-scattered radiance along a given finite path, $tau(a,b)$ denotes the transmittance between two points and $L_{inf}(observer, direction)$ denotes in-scattered radiance long a given infinite path.

A.3 Error Analysis

An error analysis has been performed for the model fitting described in chapter 4. The charts included in this appendix show every single reference image compared against a matching analytical model output.

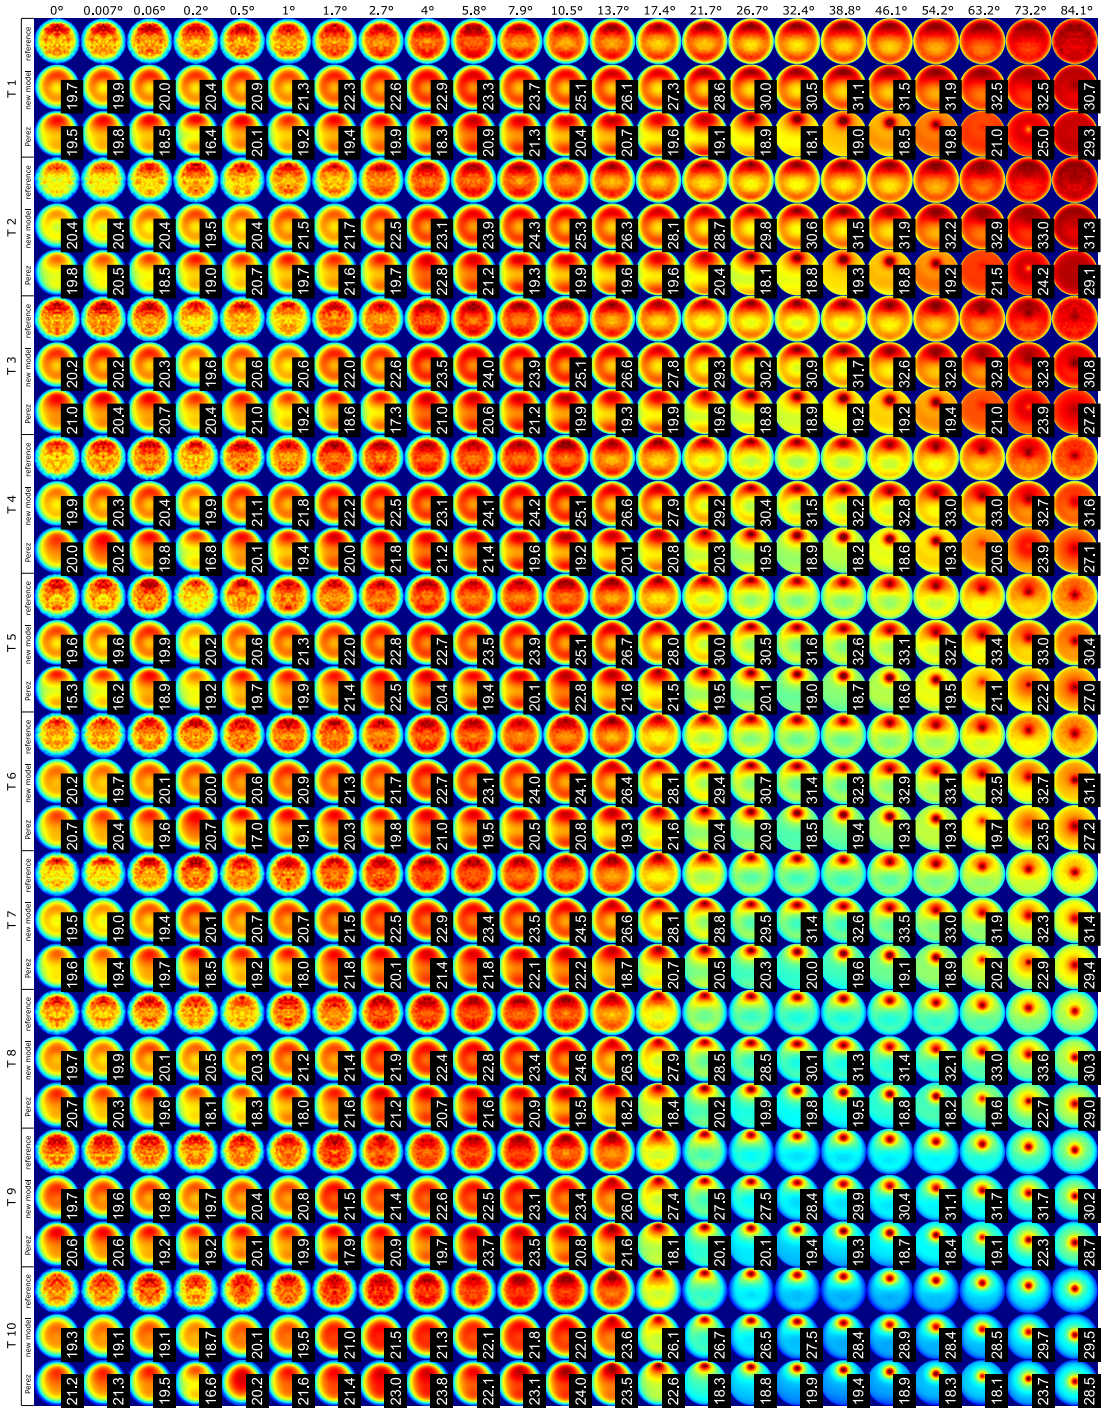

Figure A.3: False color result charts of the fitting process for $320nm$. The rows labeled "ref" are results of the brute force reference simulation, those labeled "new model" are the output of the new model proposed in the paper. Rows labeled "Perez" are the output of the original Perez formula, with distribution parameters fitted specifically separately for each combination of turbidity/elevation. T indicates the turbidity value. Logarithmic SNR measurements are shown for each model output. Note the substantially different scattering patterns compared to the visible range that are in evidence for the ultra-violet part of the spectrum – the properties of Rayleigh scattering lead to UV emission being comparatively uniformly distributed across the sky dome.

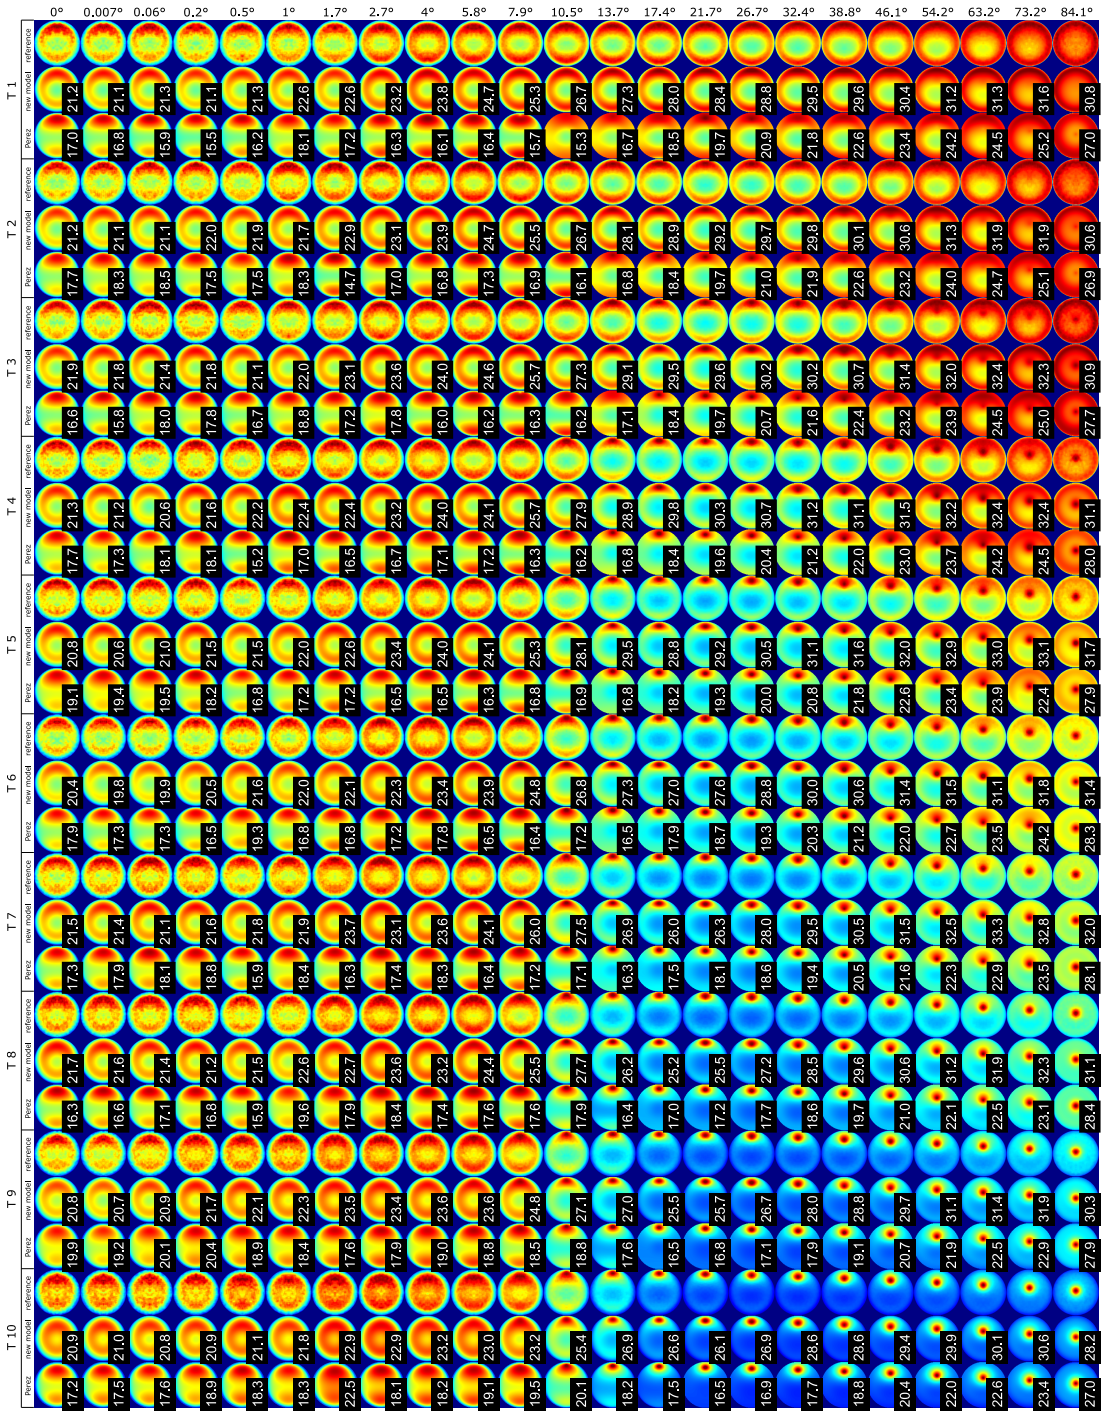

Figure A.4: False color result charts of the fitting process for $360nm$. The rows labeled "ref" are results of the brute force reference simulation, those labeled "new model" are the output of the new model proposed in the paper. Rows labeled "Perez" are the output of the original Perez formula, with distribution parameters fitted specifically separately for each combination of turbidity/elevation. T indicates the turbidity value. Logarithmic SNR measurements are shown for each model output.

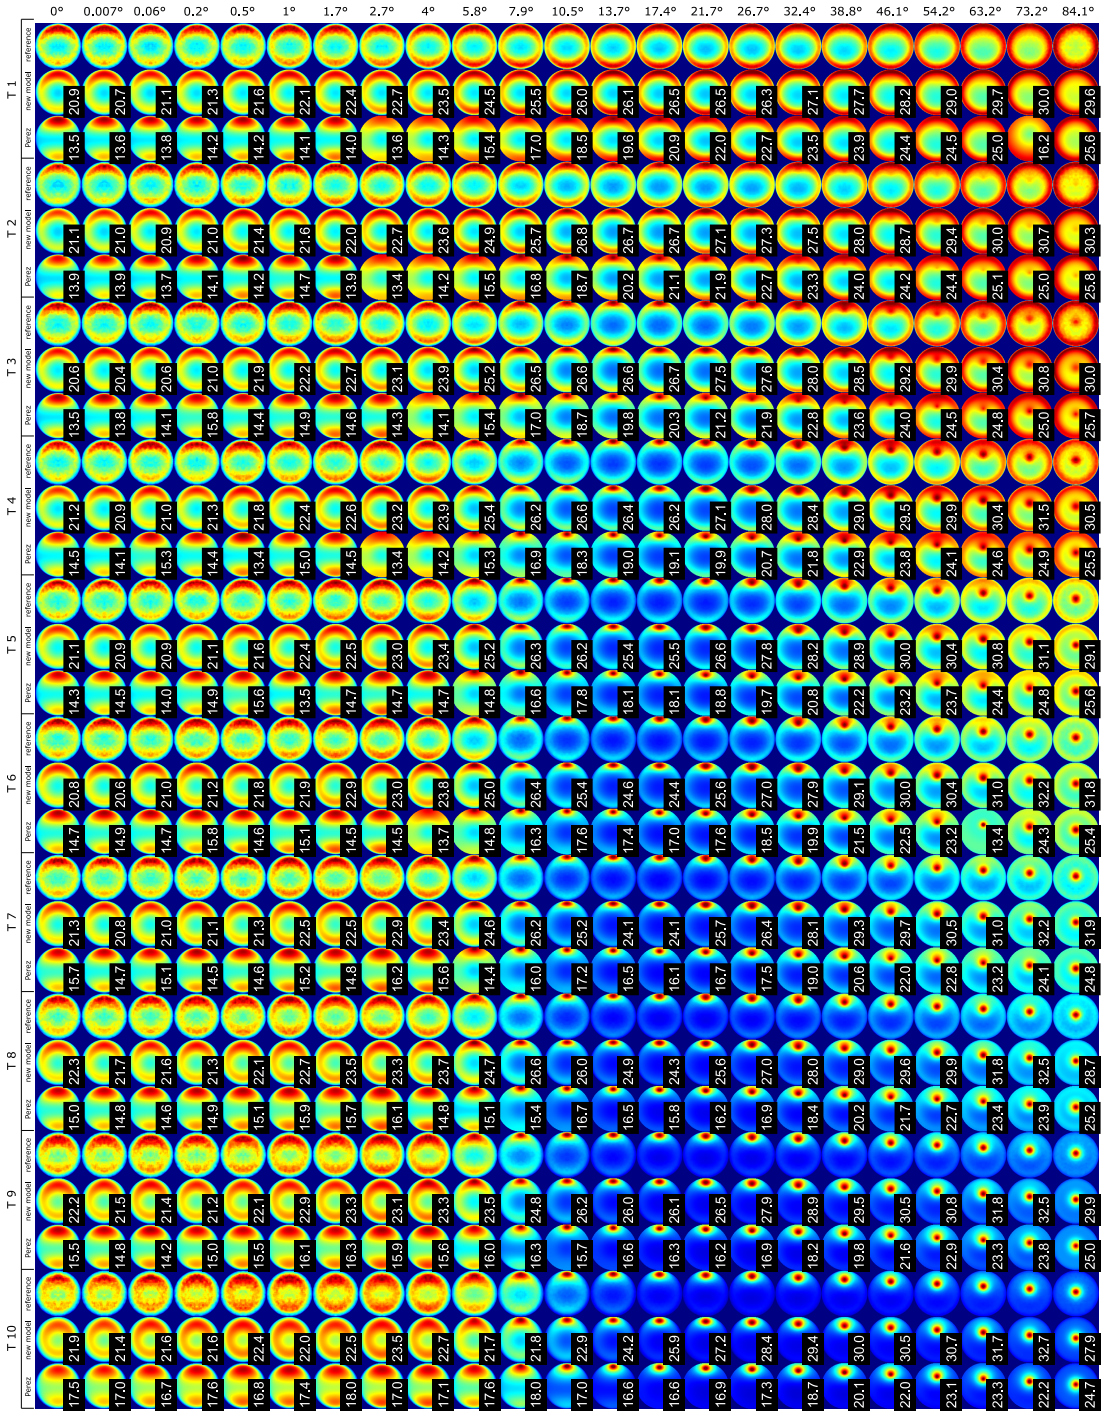

Figure A.5: False color result charts of the fitting process for $400nm$. The rows labeled "ref" are results of the brute force reference simulation, those labeled "new model" are the output of the new model proposed in the paper. Rows labeled "Perez" are the output of the original Perez formula, with distribution parameters fitted specifically separately for each combination of turbidity/elevation. T indicates the turbidity value. Logarithmic SNR measurements are shown for each model output.

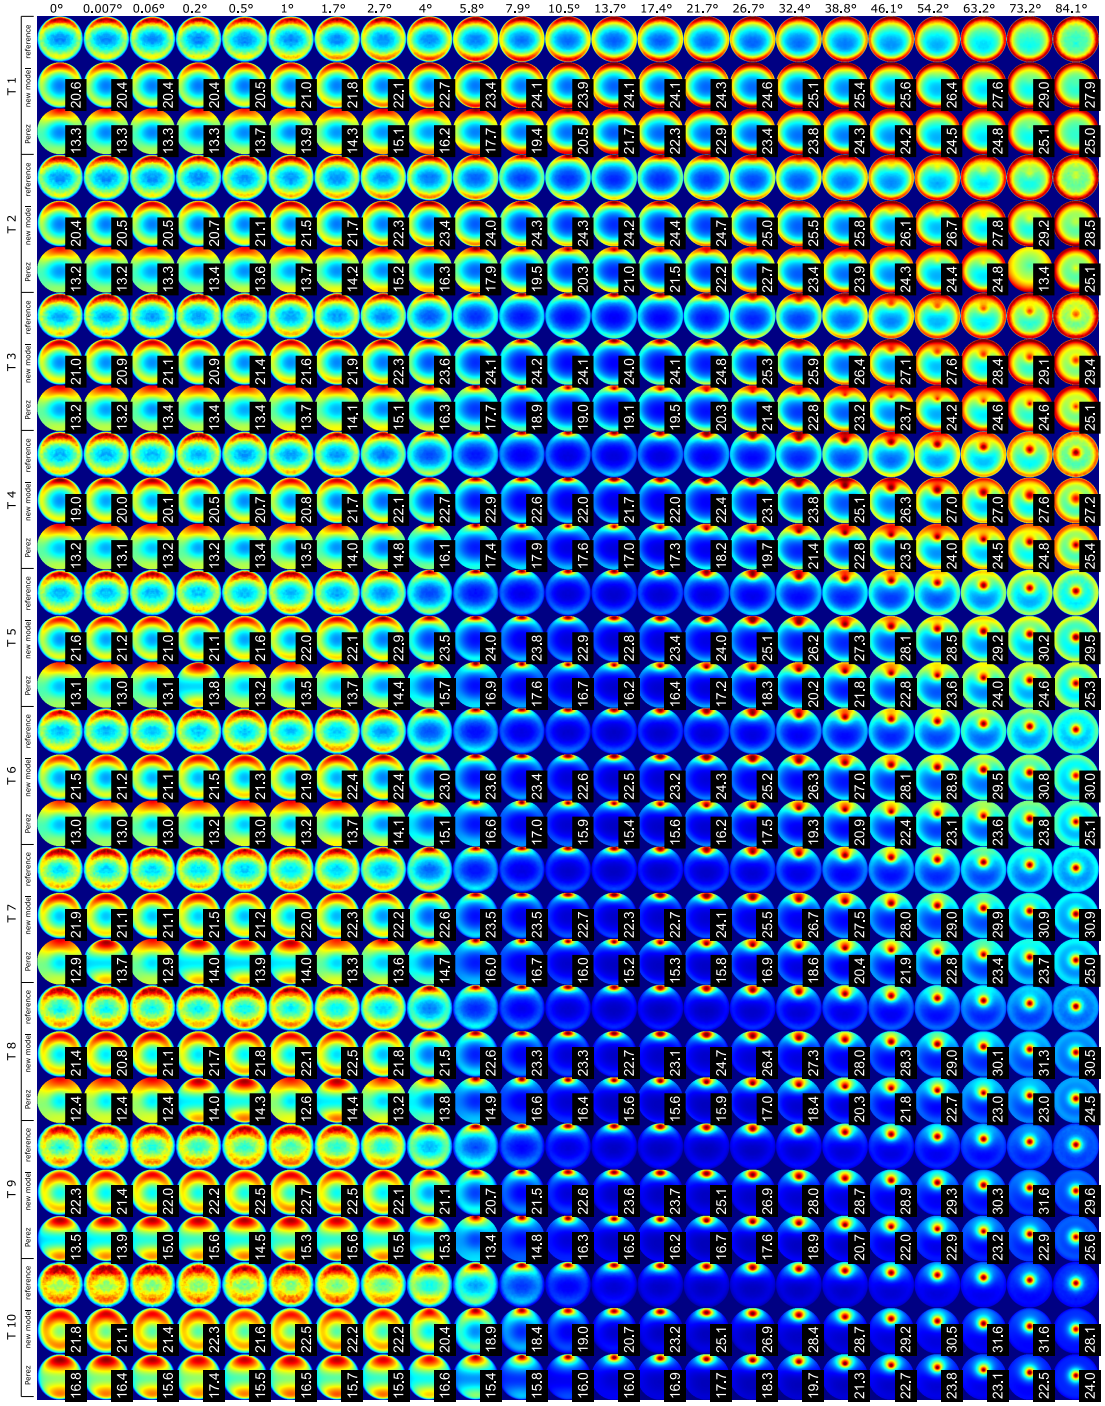

Figure A.6: False color result charts of the fitting process for $440nm$. The rows labeled "ref" are results of the brute force reference simulation, those labeled "new model" are the output of the new model proposed in the paper. Rows labeled "Perez" are the output of the original Perez formula, with distribution parameters fitted specifically separately for each combination of turbidity/elevation. T indicates the turbidity value. Logarithmic SNR measurements are shown for each model output.

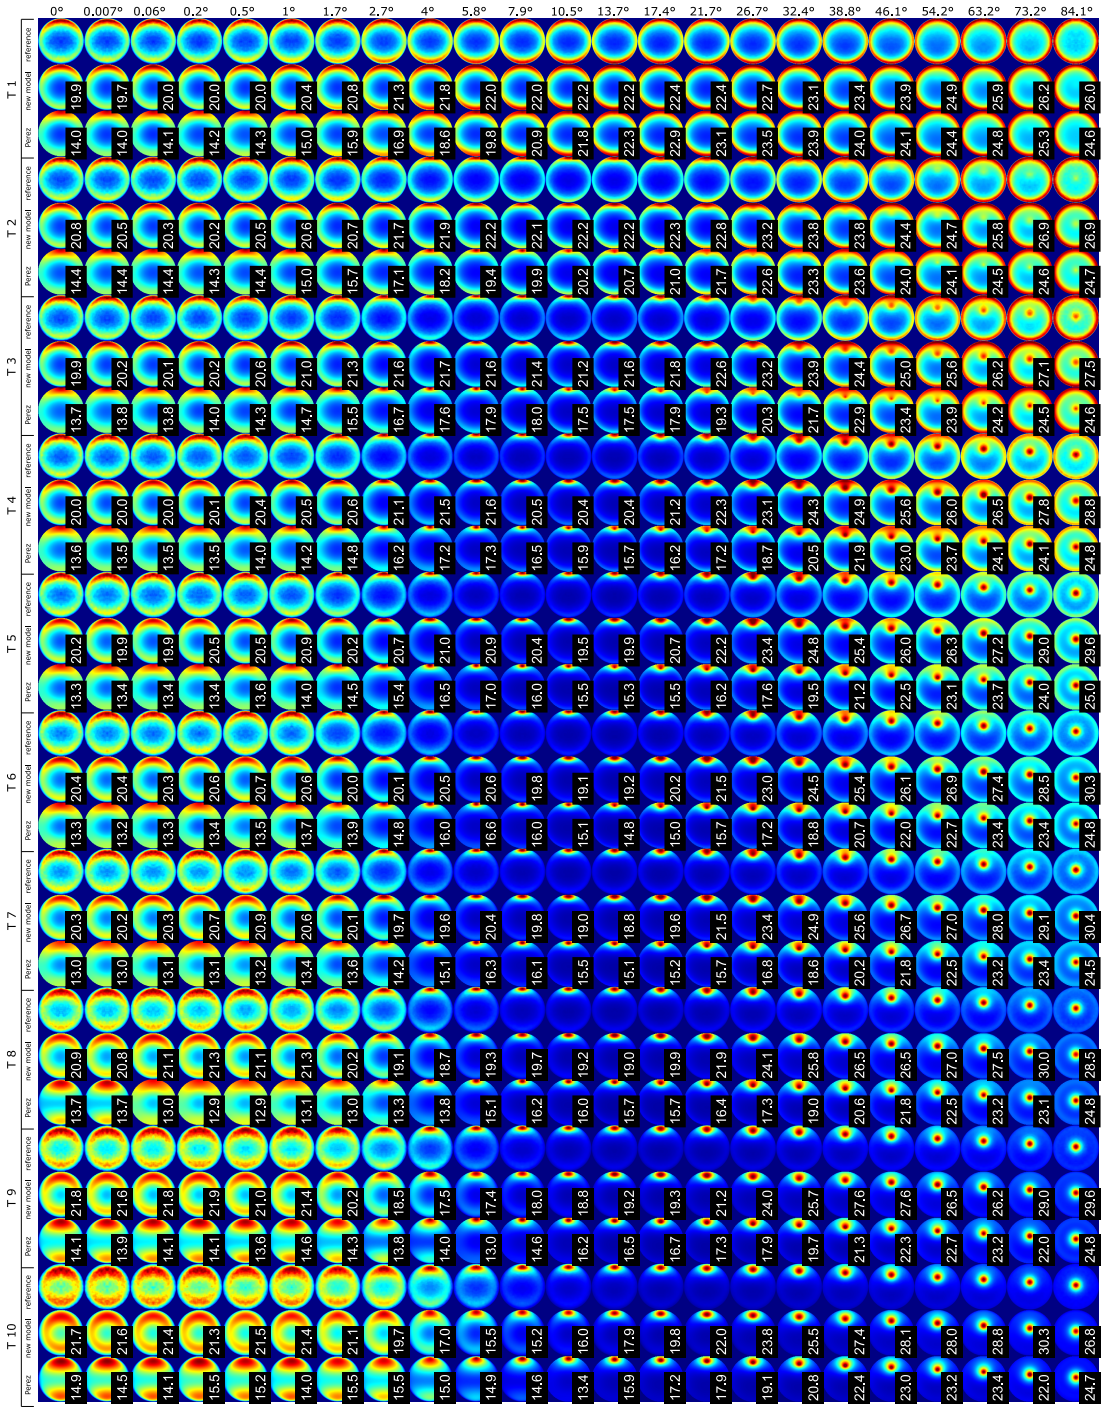

Figure A.7: False color result charts of the fitting process for $480nm$. The rows labeled "ref" are results of the brute force reference simulation, those labeled "new model" are the output of the new model proposed in the paper. Rows labeled "Perez" are the output of the original Perez formula, with distribution parameters fitted specifically separately for each combination of turbidity/elevation. T indicates the turbidity value. Logarithmic SNR measurements are shown for each model output.

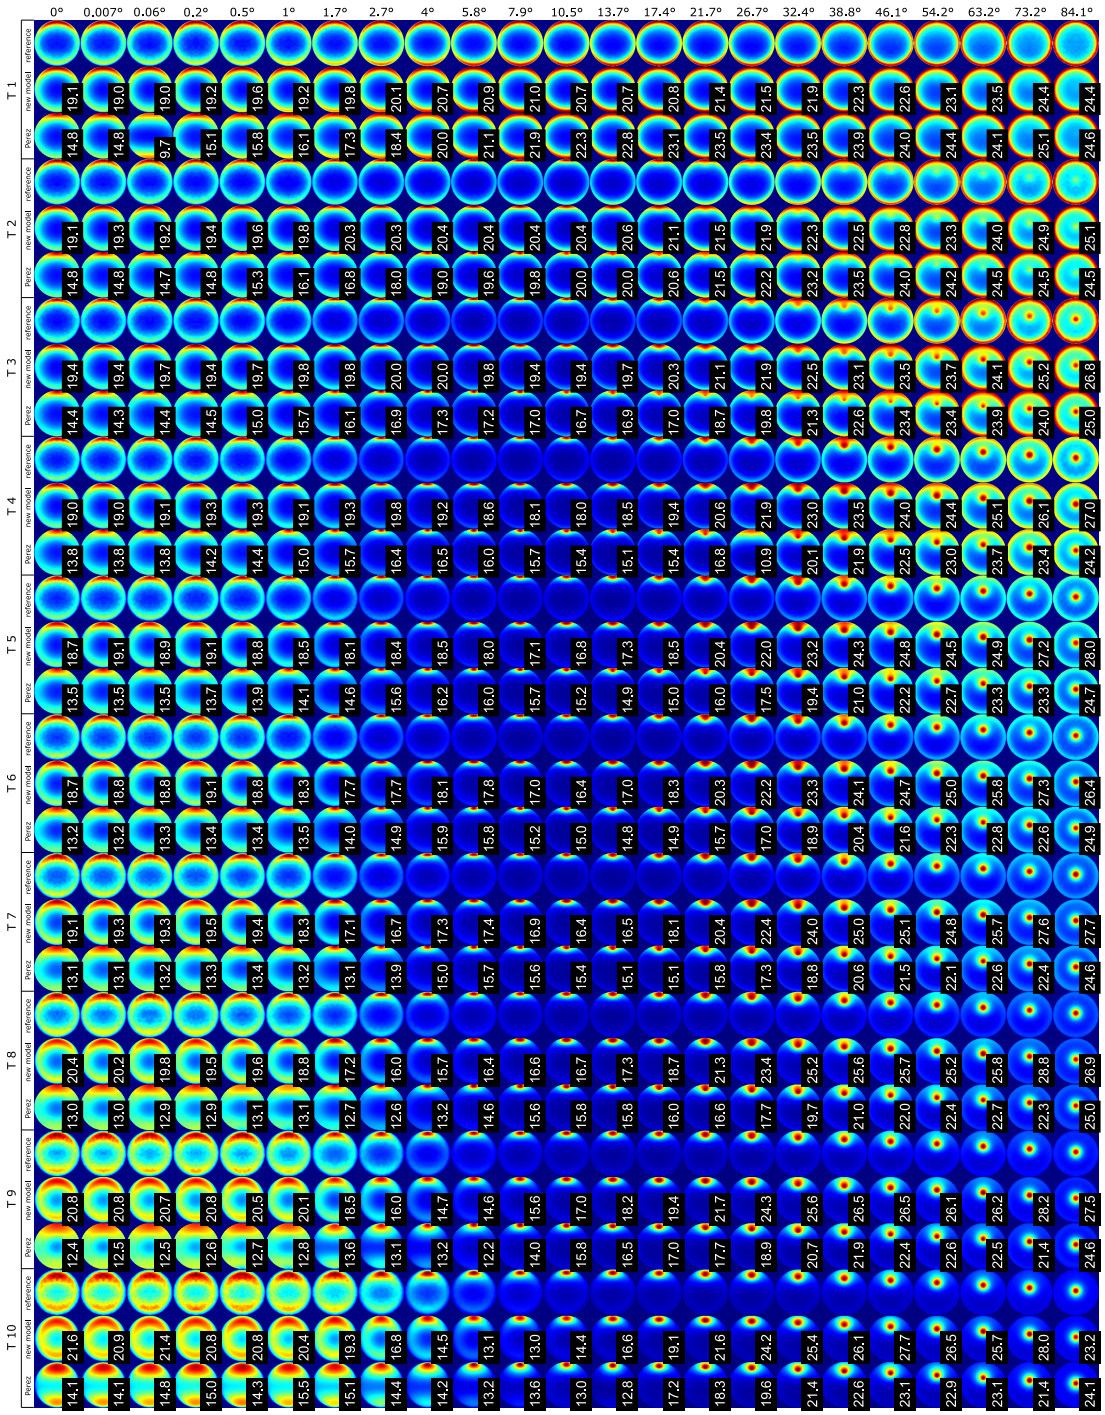

Figure A.8: False color result charts of the fitting process for $520nm$. The rows labeled "ref" are results of the brute force reference simulation, those labeled "new model" are the output of the new model proposed in the paper. Rows labeled "Perez" are the output of the original Perez formula, with distribution parameters fitted specifically separately for each combination of turbidity/elevation. T indicates the turbidity value. Logarithmic SNR measurements are shown for each model output.

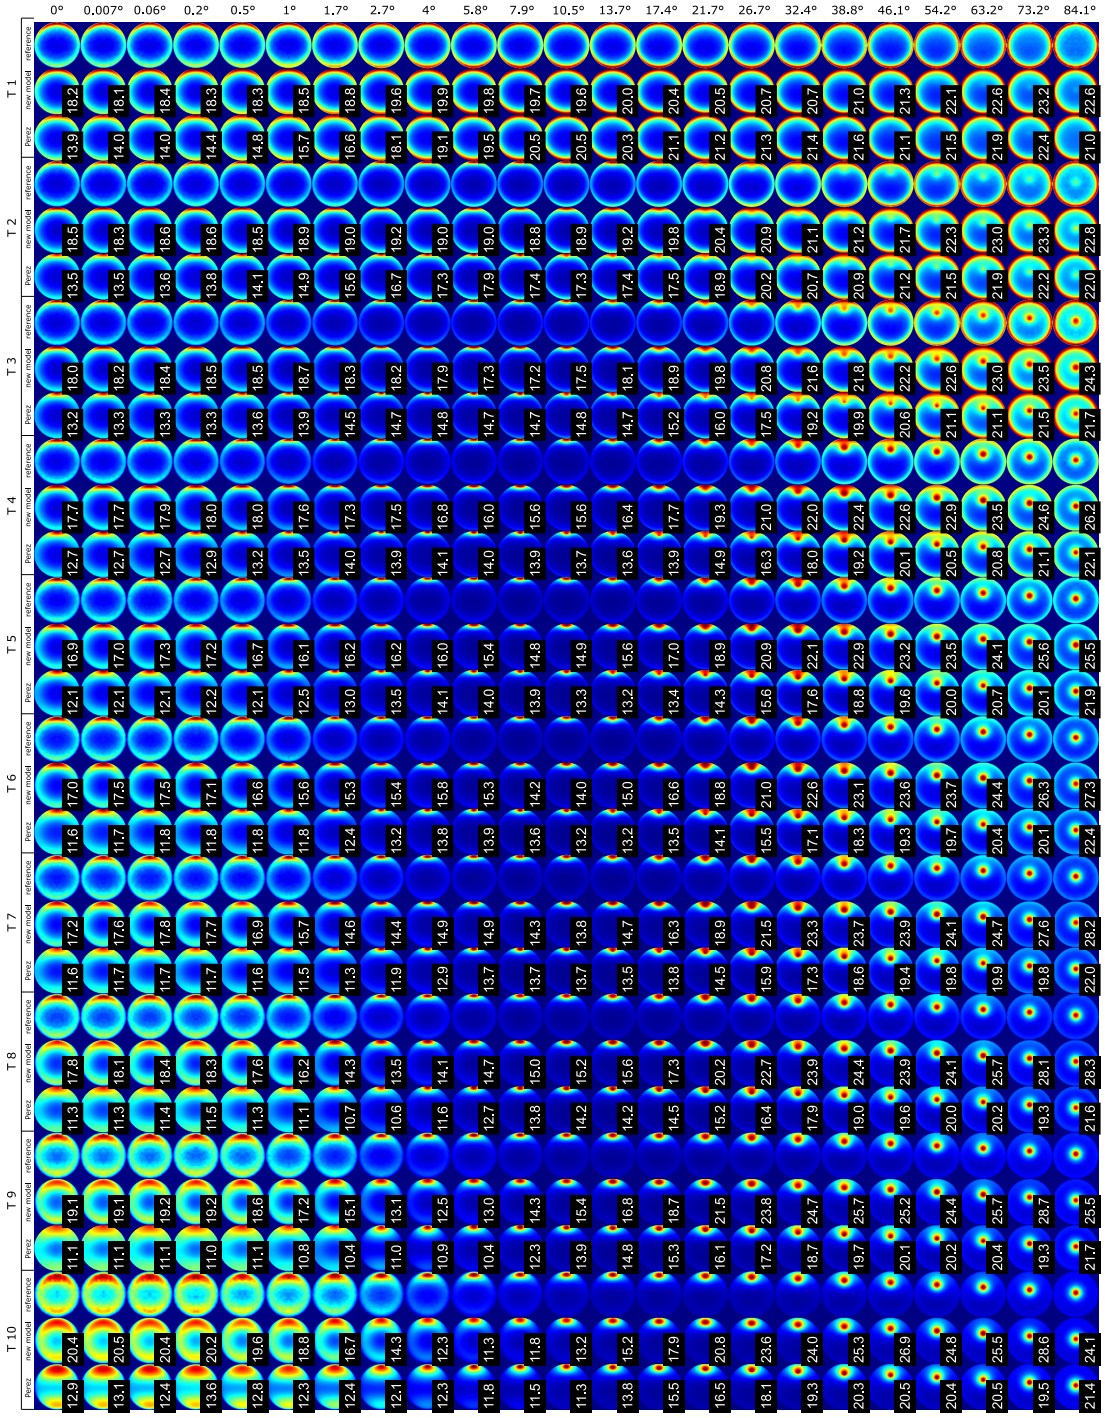

Figure A.9: False color result charts of the fitting process for $560nm$. The rows labeled "ref" are results of the brute force reference simulation, those labeled "new model" are the output of the new model proposed in the paper. Rows labeled "Perez" are the output of the original Perez formula, with distribution parameters fitted specifically separately for each combination of turbidity/elevation. T indicates the turbidity value. Logarithmic SNR measurements are shown for each model output.

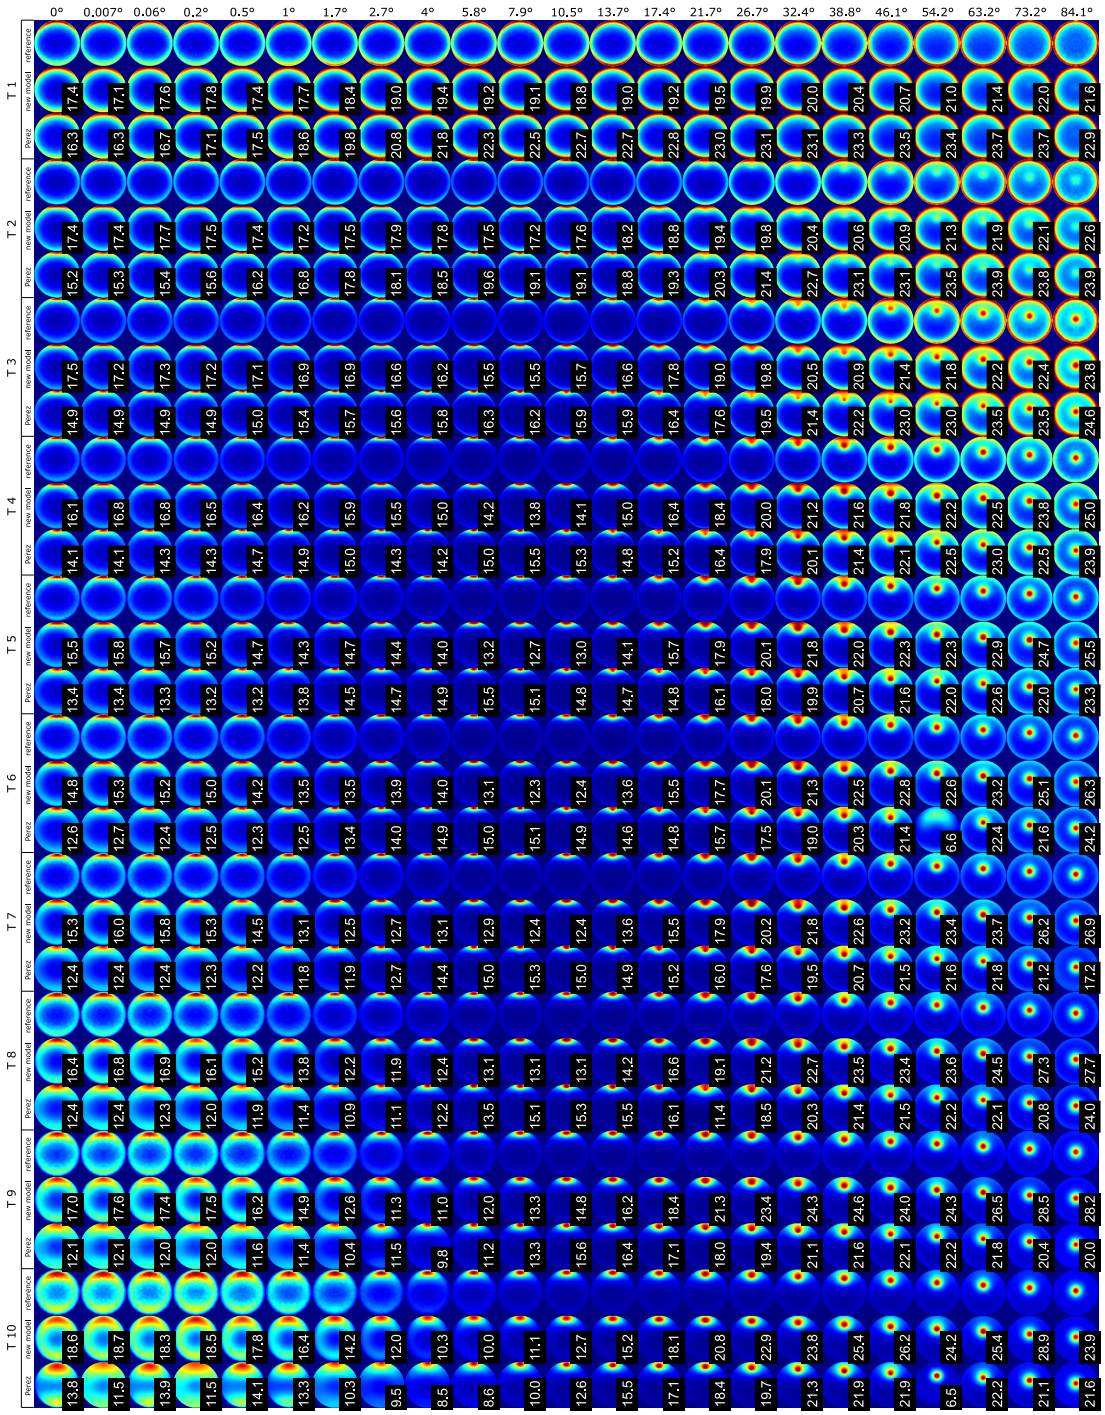

Figure A.10: False color result charts of the fitting process for $600nm$. The rows labeled "ref" are results of the brute force reference simulation, those labeled "new model" are the output of the new model proposed in the paper. Rows labeled "Perez" are the output of the original Perez formula, with distribution parameters fitted specifically separately for each combination of turbidity/elevation. T indicates the turbidity value. Logarithmic SNR measurements are shown for each model output.

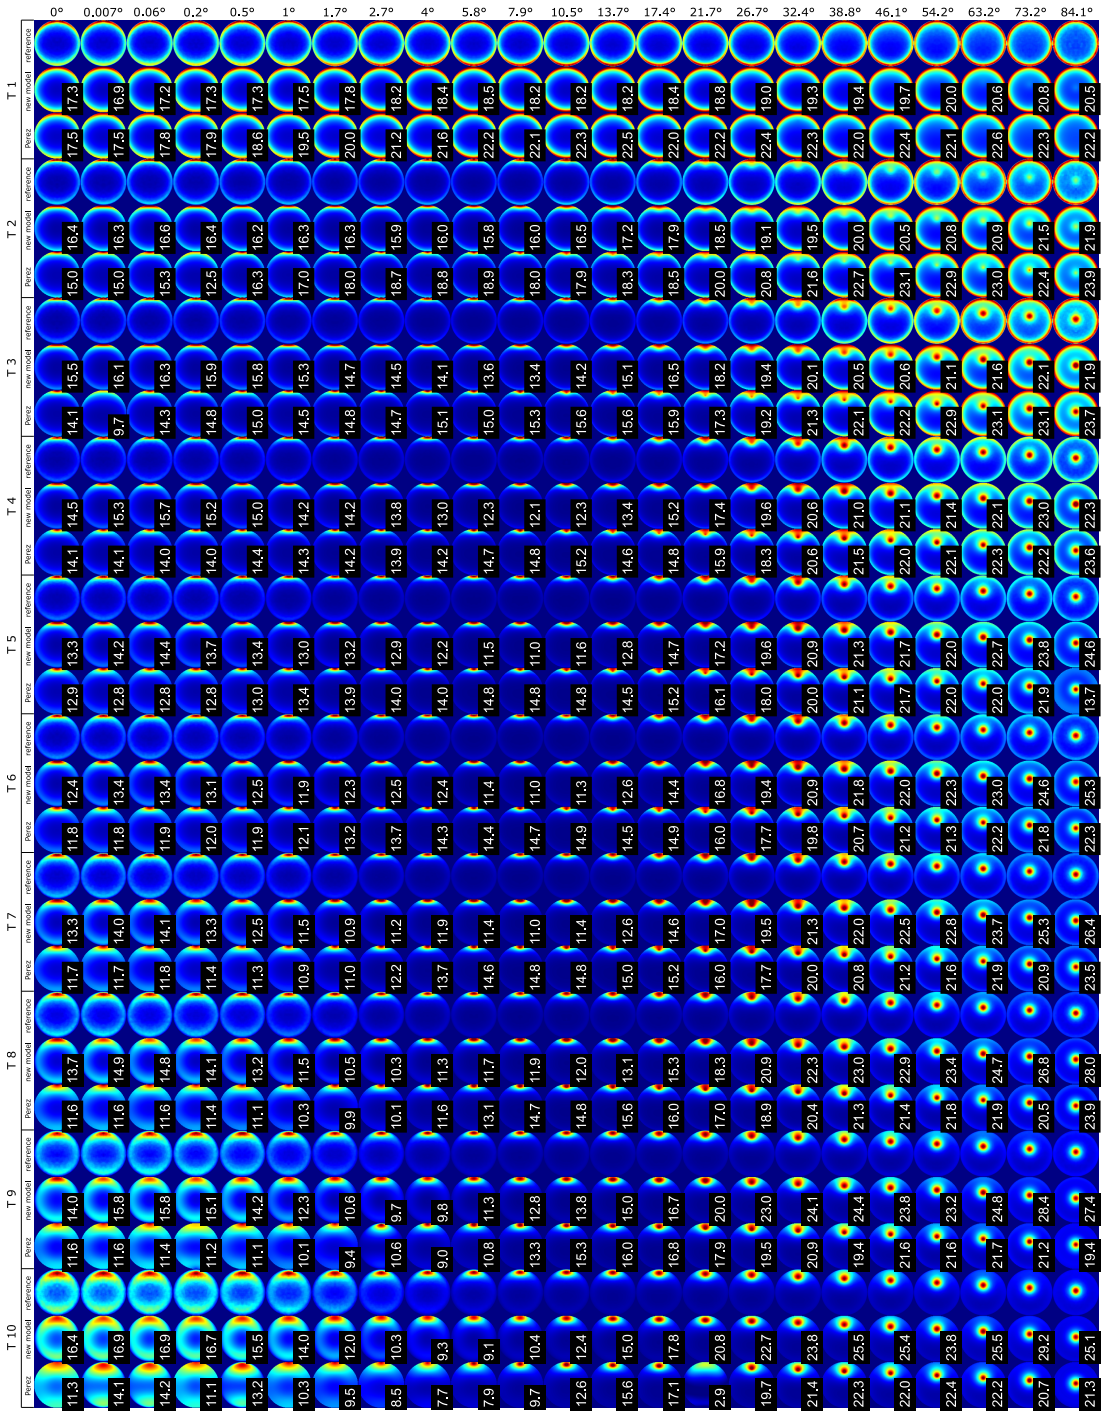

Figure A.11: False color result charts of the fitting process for $640nm$. The rows labeled "ref" are results of the brute force reference simulation, those labeled "new model" are the output of the new model proposed in the paper. Rows labeled "Perez" are the output of the original Perez formula, with distribution parameters fitted specifically separately for each combination of turbidity/elevation. T indicates the turbidity value. Logarithmic SNR measurements are shown for each model output.

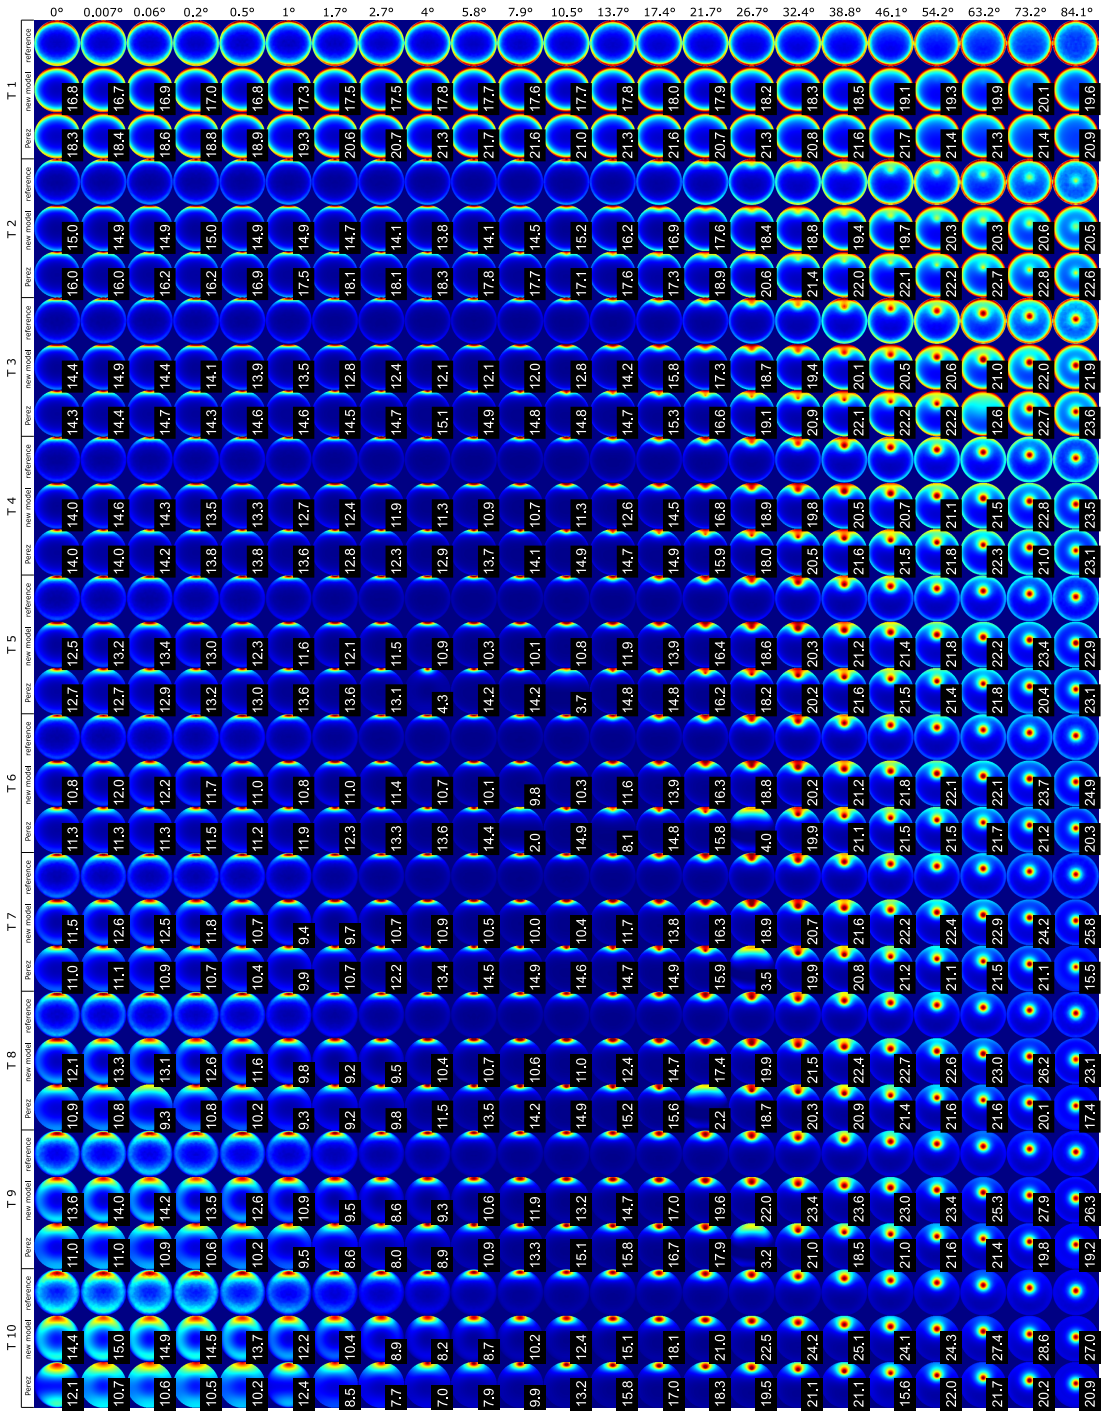

Figure A.12: False color result charts of the fitting process for $680nm$. The rows labeled "ref" are results of the brute force reference simulation, those labeled "new model" are the output of the new model proposed in the paper. Rows labeled "Perez" are the output of the original Perez formula, with distribution parameters fitted specifically separately for each combination of turbidity/elevation. T indicates the turbidity value. Logarithmic SNR measurements are shown for each model output.

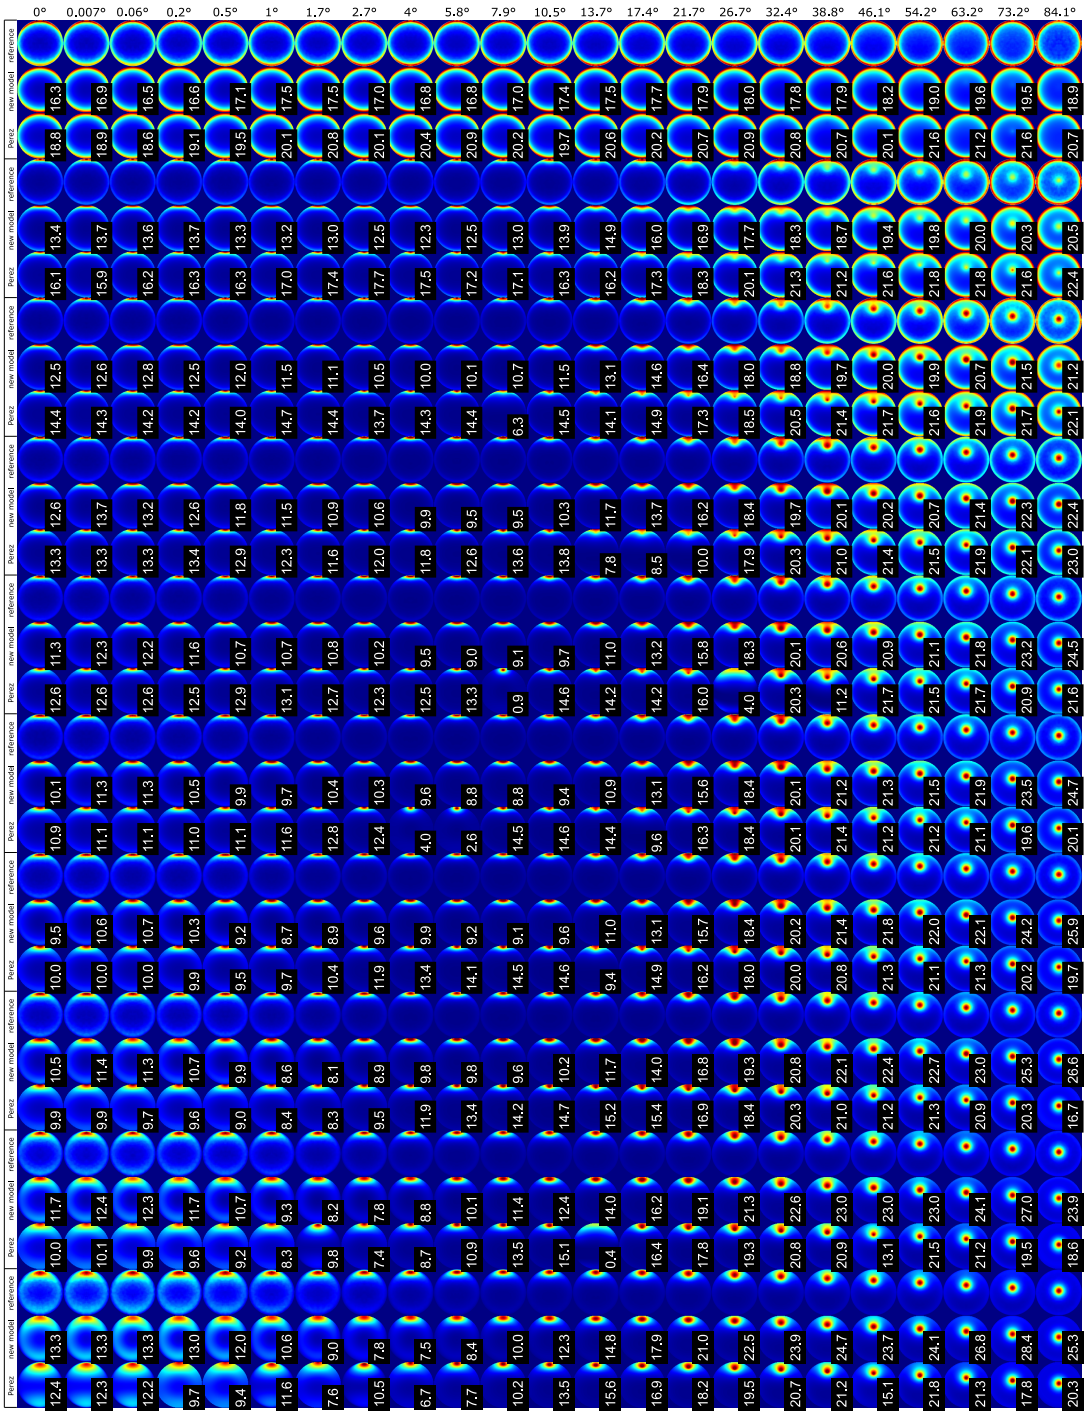

Figure A.13: False color result charts of the fitting process for $720nm$. The rows labeled "ref" are results of the brute force reference simulation, those labeled "new model" are the output of the new model proposed in the paper. Rows labeled "Perez" are the output of the original Perez formula, with distribution parameters fitted specifically separately for each combination of turbidity/elevation. T indicates the turbidity value. Logarithmic SNR measurements are shown for each model output.

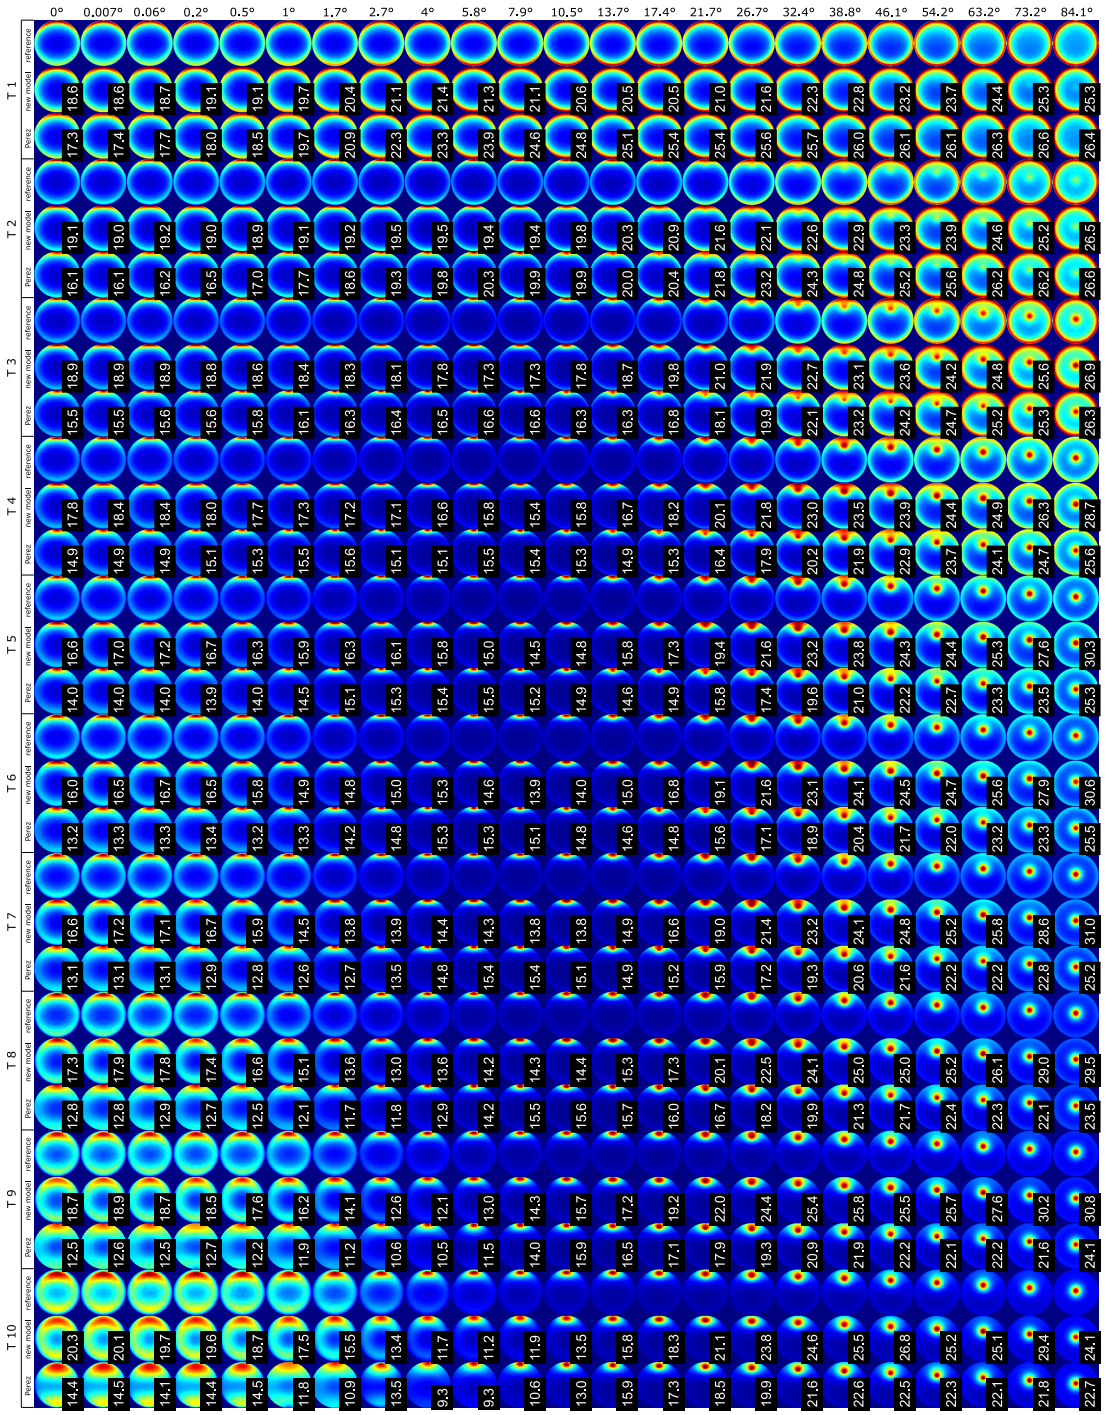

Figure A.14: False color result charts of the fitting process for the X channel of the XYZ model. The rows labeled "ref" are results of the brute force reference simulation, those labeled "new model" are the output of the new model proposed in the paper. Rows labeled "Perez" are the output of the original Perez formula, with distribution parameters fitted specifically separately for each combination of turbidity/elevation. T indicates the turbidity value. Logarithmic SNR measurements are shown for each model output.

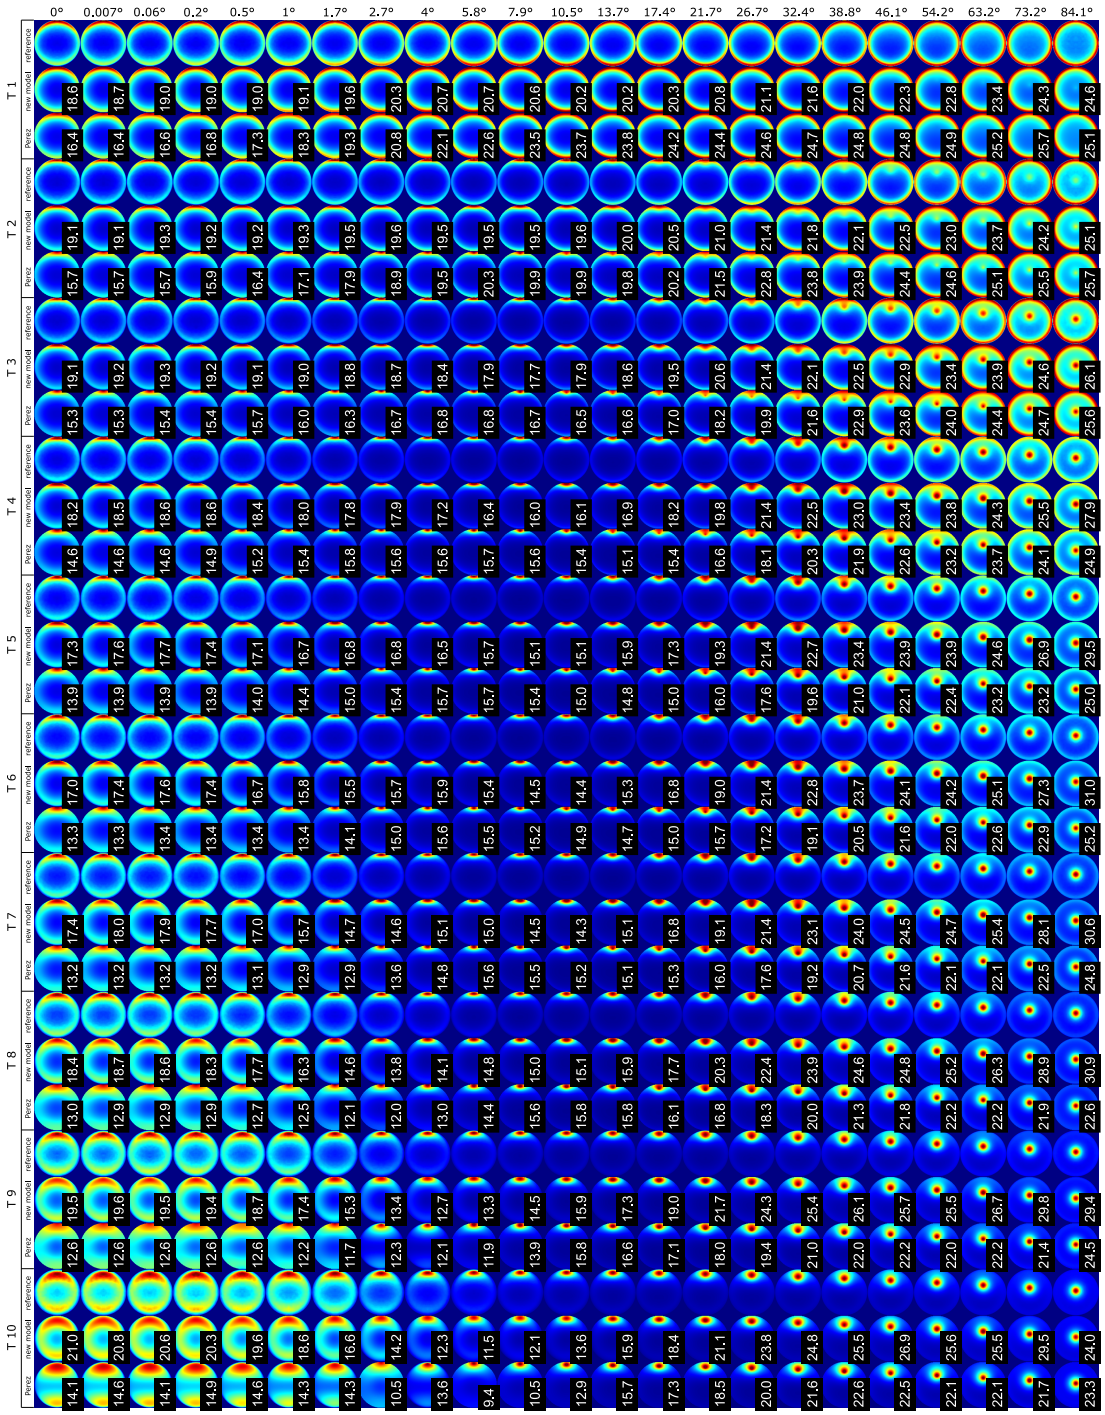

Figure A.15: False color result charts of the fitting process for the Y channel of the XYZ model. The rows labeled "ref" are results of the brute force reference simulation, those labeled "new model" are the output of the new model proposed in the paper. Rows labeled "Perez" are the output of the original Perez formula, with distribution parameters fitted specifically separately for each combination of turbidity/elevation. T indicates the turbidity value. Logarithmic SNR measurements are shown for each model output.

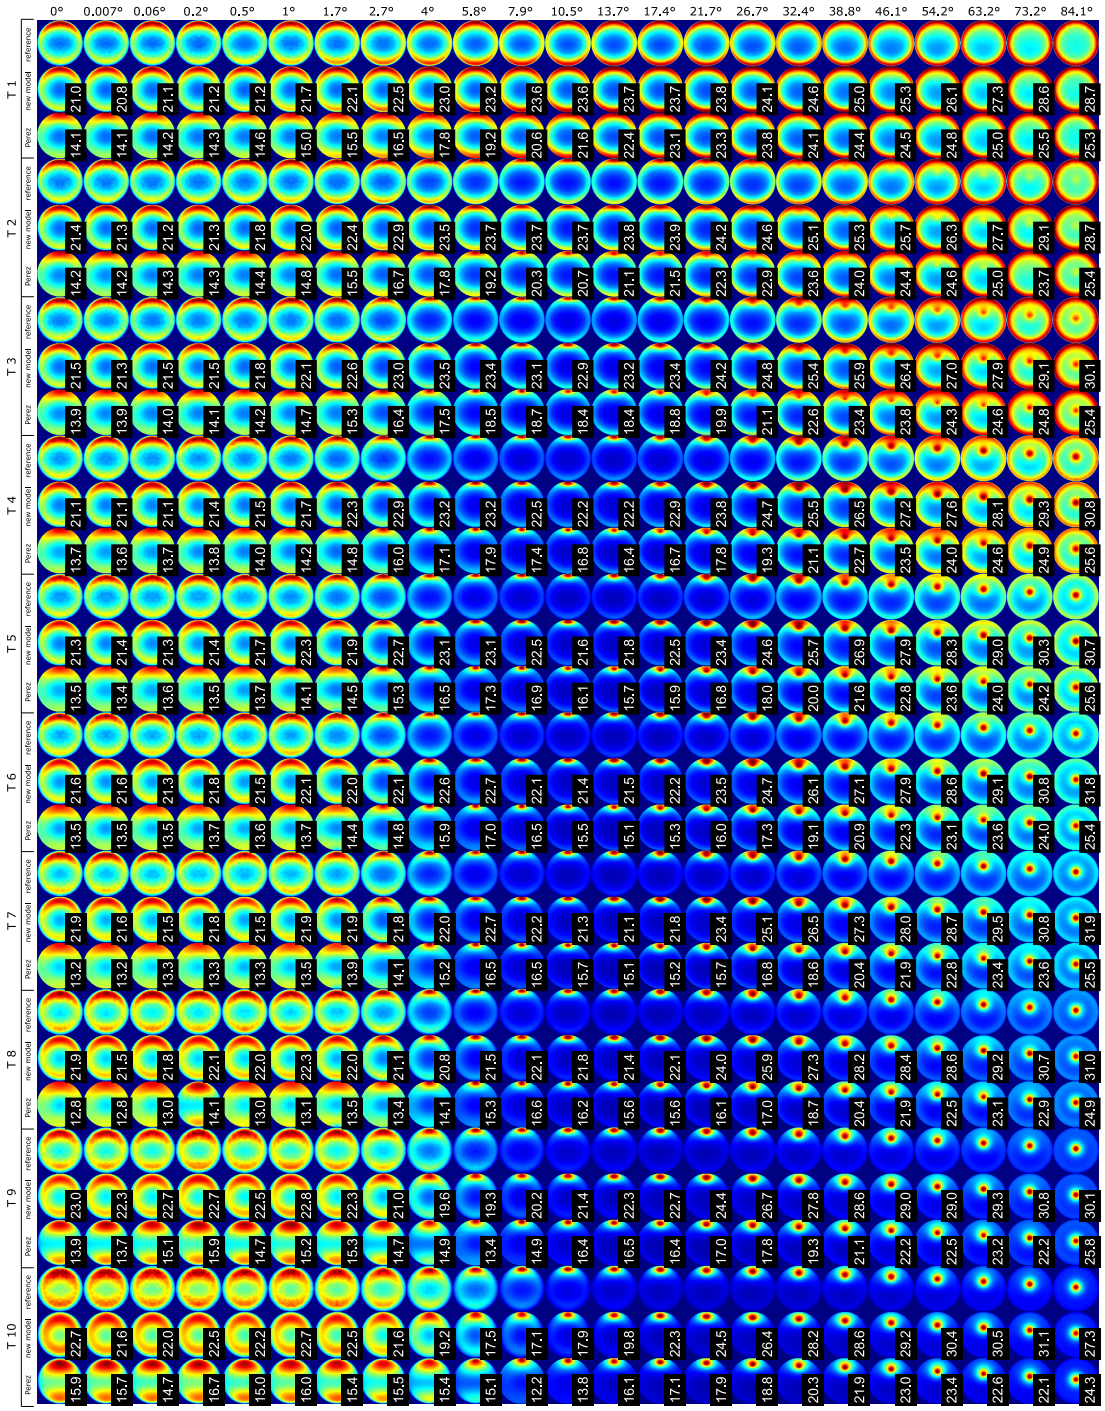

Figure A.16: False color result charts of the fitting process for the Z channel of the XYZ model. The rows labeled "ref" are results of the brute force reference simulation, those labeled "new model" are the output of the new model proposed in the paper. Rows labeled "Perez" are the output of the original Perez formula, with distribution parameters fitted specifically separately for each combination of turbidity/elevation. T indicates the turbidity value. Logarithmic SNR measurements are shown for each model output.

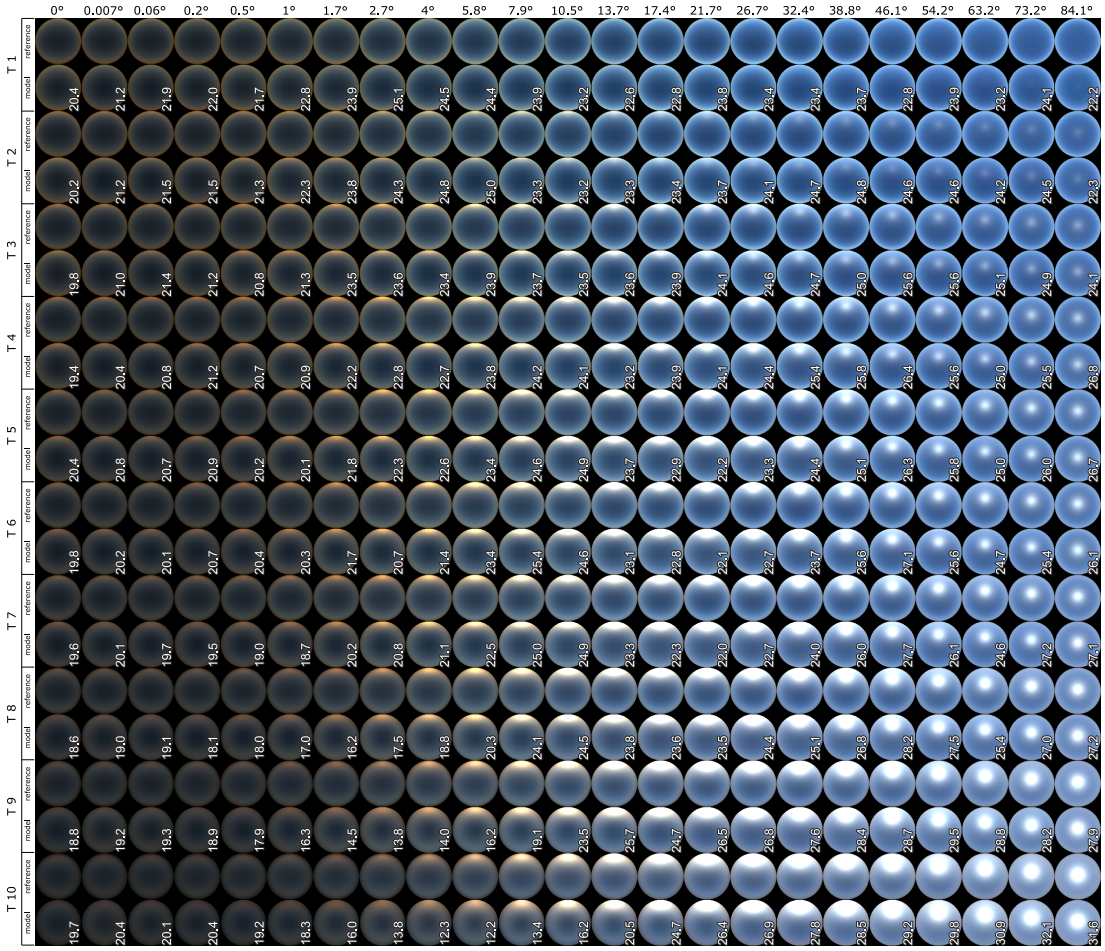

Figure A.17: Outputs of the brute force path tracer, converted to sRGB compared to outputs of the new spectral model. Logarithmic SNR measurements in Lab space are shown for each model output.