

A Fitted Radiance and Attenuation Model for Realistic Atmospheres

ALEXANDER WILKIE*, Charles University, Czech Republic

PETR VEVODA*, Charles University, Czech Republic and Chaos Czech a.s., Czech Republic

THOMAS BASHFORD-ROGERS, University of the West of England, United Kingdom

LUKÁŠ HOŠEK, Charles University, Czech Republic

TOMÁŠ ISER, Charles University, Czech Republic

MONIKA KOLÁŘOVÁ, Charles University, Czech Republic

TOBIAS RITTIG, Charles University, Czech Republic

JAROSLAV KŘIVÁNEK, Charles University, Czech Republic




Fig. 1. **Top row:** a mountain landscape rendered from 4.8 km altitude using our atmospheric model with finite-distance in-scattering. **Left:** 2° solar elevation, 30 km visibility. **Right:** 15° solar elevation, 60 km visibility. **Bottom row:** a cityscape rendered in a late afternoon setting, and post sunset.

We present a fitted model of sky dome radiance and attenuation for realistic terrestrial atmospheres. Using scatterer distribution data from atmospheric measurement data, our model considerably improves on the visual realism of existing analytical clear sky models, as well as of interactive methods that are based on approximating atmospheric light transport. We also provide

features not found in fitted models so far: radiance patterns for post-sunset conditions, in-scattered radiance and attenuation values for finite viewing distances, an observer altitude resolved model that includes downward-looking viewing directions, as well as polarisation information. We introduce a fully spherical model for in-scattered radiance that replaces the family of hemispherical functions originally introduced by Perez et al., and which was extended for several subsequent analytical models: our model relies on reference image compression via tensor decomposition instead.

*Alexander Wilkie and Petr Vevoda share the first authorship of this work.

Authors' addresses: [wilkie, vevoda]@cg.mff.cuni.cz, Univerzita Karlova, KSVI MFF, Malostranské náměstí 25, 118 00 Prague 1, Czech Republic.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

0730-0301/2021/8-ART135 \$15.00

<https://doi.org/10.1145/3450626.3459758>

CCS Concepts: • **Computing methodologies** → **Rendering**;

Additional Key Words and Phrases: sky dome models

ACM Reference Format:

Alexander Wilkie, Petr Vevoda, Thomas Bashford-Rogers, Lukáš Hošek, Tomáš Iser, Monika Kolářová, Tobias Rittig, and Jaroslav Křivánek. 2021. A Fitted Radiance and Attenuation Model for Realistic Atmospheres. *ACM Trans. Graph.* 40, 4, Article 135 (August 2021), 14 pages. <https://doi.org/10.1145/3450626.3459758>

1 INTRODUCTION

For high quality renderings of outdoor scenes, one needs realistic models of sky-dome radiance, atmospheric scattering, and optionally also cloudscapes. It has been known for a considerable time how to compute these, via brute force path tracing of realistic atmospheric and cloud models. However, the computational cost of this is still infeasible for routine production use, and will remain so for the foreseeable future – in particular for interactive use cases. Three families of techniques have established themselves to cover for this performance deficiency of full atmospheric path tracing:

- (1) *HDR sky-dome captures* are intrinsically realistic, and can include clouds: typical usage is as an HDR environment map. But they lack matching atmospheric scattering information, are static, and cannot easily be manipulated to e.g. modify solar elevation or atmospheric parameters.
- (2) *Approximative sky models* such as [Hillaire 2020] provide excellent results for interactive settings. But all techniques in this category are based on approximating light transport in the atmosphere and typically neglect higher order scattering events.
- (3) *Fitted analytical models of sky-dome radiance* that are based on brute force simulations of atmospheric light transport have proven popular for use in offline rendering. Due to being based on physical simulations, such models can, at least in theory, deliver visual fidelity on par with HDR captures: but with the added ability to modify sky appearance and solar position. And crucially, such models can also use the non-trivial vertical scatterer profiles found in real atmospheres.

In spite of the impressive performance of current interactive techniques, they cannot be used if high degrees of accuracy are required, due to their use of simplified light transport, and due to performing computations in a colour space instead of using spectral rendering.

For several application areas, such as movie VFX, or predictive rendering applications like training of autonomous vehicle sensor suite software, one needs more realistic models: and especially for the latter, reliable spectral data for a wide range of configurations is also needed. As brute-force approaches are too slow, fitted models will remain in use: but improvements in this area are needed, as existing techniques are all lacking a number of crucial features.

In particular, altitude-dependent in-scattering and absorption data that accurately matches the in-scattered light component for observer altitudes above ground level are missing from current fitted models. Both these features are crucial for the ability 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 running an expensive full atmospheric scattering simulation. Also, for accurate simulations of specular object appearance in outdoor scenes (again a feature needed for autonomous vehicle sensor training), it can be important to consider the polarisation of sky-dome radiance: and no such functionality is available in existing models.

In this paper, we advance the state of the art with regard to the third kind of approach. We propose an integrated model of sky-dome radiance and attenuation, and follow the general approach of previous models of the third type by first running brute force

simulations, and then fitting a model to the obtained data. But we advance the state of the art with regard to practically all components of this process. The main contributions are:

- The use of reference data from atmospheric science to define realistic vertical scatterer distribution profiles: these profiles are then used in a polarisation-aware path tracer to generate a large database of polarised reference images.
- Fully spherical reference images, which are generated for a range of observer altitudes up to 15 km: this is a considerable improvement over current hemispherical models that are only defined for ground-based observers.
- Verification of these reference images against the output of dedicated atmospheric simulation software.
- A new tensor decomposition approach to compress the reference image dataset. With it, artefact-free interpolated sky-dome images can be reconstructed from a coefficient set that is a fraction of the size of the reference images themselves.
- Solar elevations up to -4.2° are included in the model, as the new compression technique is powerful enough to deal with the changing sky dome luminance patterns after sunset.
- A matching model for atmospheric transmittance is provided.
- And finally, there is also a matching model for sky-dome polarisation.

In the remainder of this paper, we first review the state of the art in such models within the field of computer graphics in Sec. 2. We then discuss the physics background of sky-dome luminance and attenuation patterns in Sec. 3, and our atmospheric model and our brute force path tracer are discussed in Sec. 4. The fitting process itself is described in Sec. 5, and we conclude by demonstrating the results of our new model in Sec. 6.

2 STATE OF THE ART IN SKY MODELS

As already outlined in the introduction, work on sky dome radiance and atmospheric transmission falls into four broad categories within computer graphics:

- (0) brute-force simulations of light transport in the atmosphere
- (1) capture and measurement of real skies
- (2) interactive approximations to atmospheric scattering
- (3) fitted analytical models

True brute force approaches (here, numbered zero for consistency with the introduction) are not nearly fast enough for production use, so the remaining three can all be seen as means to provide adequate levels of realism at acceptable performance levels. For brute-force simulations and fitted analytical models there exists a considerable body of work within the atmospheric research community, for which we provide a summary in Table 1, and we refer the reader to [Bruneton 2016] for a comprehensive overview and evaluation of fitted analytical models and brute-force solvers. In the remainder of this section we provide more details about selected work from each category.

2.1 Brute-force Simulations

Examples of this category, such as the work of Nishita et al. [1996; 1993], Haber et al. [2005], Bruneton et al. [2008] and Guimera et al. [2018], tend to yield excellent results in terms of visual quality.

Most of them actually contain a pre-computation step that reduces the complexity of the models' evaluation at the time of rendering, albeit for the set of parameters (observer altitude, sun elevation) fixed at the pre-computation step. Real-time methods, such as O'Neal [2005] compute approximations to atmospheric lighting; these typically trade effects such as multiple scattering for fast computation. 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. These can serve as a valuable source of reference solutions for graphics research, such as those shown in the work of Wang et al. [2016]. However, for the reasons we discuss in Sec. 4.4, and notwithstanding all its excellent capabilities, libradtran is not well suited for general rendering tasks, and would not have been a good solution for the very specific problem of obtaining reference images for our own model.

2.2 Fitted Sky Dome Models

This category contains a number of models that have seen fairly widespread industrial use [Hošek and Wilkie 2012, 2013; Perez et al. 1993; Preetham et al. 1999]. These models are all attempts to fit parametric functions to the properties observed in brute force sky dome radiance simulations: common to all is their usefulness for fast, high-quality scene modelling and rendering. The main drawbacks of these models are typically a limited parameter range, a lack of realism for some regions of their parameter space, and a comparatively narrow applicability: all these models assume the observer to be at ground level, a shortcoming which we are explicitly addressing in the proposed new model. [Hošek and Wilkie 2012, 2013] also lack a dedicated atmospheric transmission model that matches the sky dome luminance. Our work adds these features, albeit at the cost of requiring precomputation for the required range of parameters, and resulting in a model which has larger memory requirements than previous techniques. Fig. 2 illustrates the conceptual difference of these models to a full solution like our model provides, while Fig. 3 shows the consequences of omitting individual components.

2.2.1 Sky Dome Polarisation Patterns. An addition to the line-up of analytical models was the polarisation work of [Wilkie et al. 2004], where they proposed an approximate analytical fit of sky dome polarisation patterns. As the authors neither had access to reference data, nor to a brute force simulator capable of including the phenomenon, they (as per their own admission in the paper) entirely based their approach on indirect reasoning. And as [Wang et al. 2016] conclusively showed, their results were rather sub-optimal. The authors of the latter paper also proposed a new, high-quality analytical model of sky dome polarisation that can be efficiently evaluated. Without an accompanying analytical model of sky dome radiance and transmission, however, the utility of this promising new model is of limited use for producing highly accurate renderings. By contrast, our new proposed model provides an integrated solution that covers all aspects of sky dome radiance (in-scattered radiance, attenuation, polarisation).




Fig. 2. Capabilities of pre-computed sky radiance models. Existing models such as [Hošek and Wilkie 2012] correspond to case A: they are quite simple, and only provide information for paths that directly go to space without hitting any objects, and with the observer at ground level. They also do not provide a model for attenuation over finite viewing distances. But for non-trivial renderings, one *additionally* needs to cover case B: with an observer viewpoint that is not necessarily on the ground, with in-scattered radiance information for finite viewing distances when objects are hit, and with matching attenuation information. Our model provides all this.

2.3 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, and collections of representative stock HDR sky images are a common lighting resource in production environments. 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




Fig. 3. Consequences of omitting some of the pre-computation components identified in Fig. 2. **Top:** a full rendering, observer altitude 8km above ground, close to sunset. **Bottom left:** the [Hošek and Wilkie 2012] model neither has a concept of observer viewpoints above ground level (so downward looking rays need to extrapolate, and solar radiance stays the same at all altitudes), nor provides an expression for in-scattered radiance or attenuation for finite distances. **Bottom right:** our model with in-scattered radiance for finite viewing distances switched off: as one can see, this component is absolutely crucial for outdoor scene realism.

	sun below horizon	arbitrary observer altitude	spectral	polarisation	approach b (brute-force) f (fitted model)
Nishita93	+	+	-	-	b
Nishita96	+	+	-	-	b
Preetham	-	-	-	-	f
Haber	+	-	-	-	b
Bruneton	+	+	+	-	b
Elek	+	+	+	-	b
Hosek	-	-	+	-	f
libradtran	+	+	+	+	b
Ours	+	+	+	+	f

Table 1. Comparison of several clear sky models and their features.

for renderings done in absolute units, or as basis for verification of fitted sky-dome models.

A 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 does not contain the exact atmospheric parameters at the time of capture, but only meteorological data. This lack of scatterer profile data means that it is not easily possible to do a brute force rendering 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. However, as the comparison in Figure 7 shows, the U.S. Standard Atmosphere, which we use as basis of our own model, is a reasonable fit for the measurements.

2.4 Transmission Models

Models of atmospheric transmission have been provided in some sky dome research papers, such as [Preetham et al. 1999], however these typically rely on simplifications such as exponential distributions of aerosols. Other approaches for transmittance such as [Hansard 2019] operate in image space, so cannot be easily applied to multiple bounces of lighting. Our work provides a specialised transmittance parametrisation and fit, is accurate to the underlying atmospheric configuration, and can be used when computing indirect lighting.

3 PHYSICS BACKGROUND FOR ATMOSPHERIC RENDERING

In one way or another, whenever one wishes to generate images of outdoor scenes, some kind of atmospheric rendering is done. As outlined in the previous section, graphics usually uses approximate solutions, to achieve reasonable performance. In this section, we briefly review physics aspects of the full problem.

A complete description of light transport in participating media, such as in atmospheres, is given by the radiative transfer equation [Subrahmanyan 1960], for which an exhaustive introduction is given by Pharr and Humphreys [2010]. It is subject to the following parameters:

- (1) The *transport coefficient* or *extinction coefficient* σ_t [m^{-1}], which is the inverse of the *mean free path*. This parameter can be obtained by multiplying the *extinction cross section*

[m^2] by the *particle concentration* [m^{-3}], i.e., the number of particles per cubic meter.

- (2) The *single scattering albedo* α (unitless), which gives the probability that a collision event in the medium is a scattering event as opposed to absorption.
- (3) The *phase function* and its parameters, which describes the directional distribution of light after a scattering event.

3.1 Scattering Mechanisms

In an atmosphere, two main scattering mechanisms prevail: Rayleigh and Mie scattering. Both can be characterised by their respective phase functions and extinction cross sections. Note that other types of mechanisms, e.g. scattering by non-spherical particles such as ice crystals, are not present in all atmospheric configurations, and we do not take them into account in our model.

Rayleigh scattering describes the interaction of light with gas-phase constituents of the atmosphere such as N_2 and O_2 . Closed-form expressions are available for calculating the cross section and phase function [Bodhaine et al. 1999; Mätzler 2002]. The particle concentrations change with altitude – we describe the vertical profiles later in Sec. 3.2.

Mie scattering describes the interaction of light with aerosols, or more generally scatterers that are larger than the light wavelength. In an atmosphere, the particles have various size distributions [Hess et al. 1998], which influences the cross sections and phase functions. As the scattering favours forward directions, it produces characteristic coronas around light sources in foggy environments.

Unlike Rayleigh scattering, simple closed-form formulas are not available for Mie scattering, and it has to be approximated [van de Hulst 1957] or pre-computed from the size distributions and tabulated separately for each wavelength [Emde et al. 2016]. In our model (Sec. 4.2), we use pre-computed extinction coefficients, single scattering albedos and asymmetry parameters together with a closed-form Henyey-Greenstein phase function [Henyey and Greenstein 1941], the validity of which we discuss in more detail in our supplementary PDF.

Both Rayleigh and Mie scattering are polarising light-matter interactions: but for sky dome scenes, the macroscopically resulting polarisation is often rather weak in the case of Mie scattering.

3.2 Actual Composition of a Typical Atmosphere

Physically correct simulation of light in an atmosphere requires understanding the atmosphere composition. In physical simulations, an atmosphere is considered to consist of molecules, aerosols and clouds, where clouds can be further divided into water clouds and ice clouds [Emde et al. 2016]. Since our model only deals with clear skies, we do not discuss clouds further.

3.2.1 Gas-phase Constituents. For the gas-phase constituents, reference atmospheres are available such as the U.S. Standard Atmosphere [Anderson et al. 1986] with tabulated concentrations in altitudes ranging from 0 km up to 120 km (Fig. 4, first and third plot). The two main constituents of air, N_2 and O_2 , form roughly 78.1 % and 20.9 % of air, respectively, up to the altitude of about 90 km, where the overall air concentration is already negligible.




Fig. 4. Selected optical properties of our atmosphere. When multiple curves are displayed, our fitted model uses the one with a solid blue line. **From left to right:** Vertical air concentration profile in a U.S. Standard Atmosphere (data: libradtran [Anderson et al. 1986; Emde et al. 2016]). O_3 absorption cross sections at the temperatures 223 K and 293 K, our model uses 223 K (data: Gorshelev et al. [2014]). Vertical O_3 concentration profiles from various sources, our model uses U.S. Standard Atmosphere again, but we also include measurements from Kaifol et al. [2012] for normal and ozone hole conditions to show that atmosphere profiles are not constant and change over time and place. Finally, the vertical aerosol concentration profiles for the viewing distances covered by our model. Here, we only show water-soluble (WASO) particles for clarity, the other particle categories are qualitatively similar. Note that the solid lines represent concentrations which are slightly smoothed versions of the OPAC dataset atmospheres (Fig. 4, right), while the dashed lines are inter- resp. extrapolated. Please consult our supplementary PDF for validations and details.

These molecules are responsible for Rayleigh scattering and are mainly responsible for the blue colour of the sky. However, as we show in Fig. 5, there is another molecule which is relevant to sky-dome appearance: ozone (O_3 , Fig. 4). Its importance for sky dome appearance in dusk conditions has long been known in atmospheric research literature [Hulburt 1953], although a recent publication has called the *direct* correspondence between twilight sky colour and ozone concentrations into question [Lee et al. 2011].

But even if there is no direct link, if skies with low solar elevations are to be rendered correctly, the inclusion of O_3 is a necessity, so we included it in our model (Sec. 4.2). So far, it has mostly been omitted in computer graphics sky models [Nishita et al. 1993; Preetham et al. 1999], and even though two models already include it [Haber et al. 2005; Kutz 2013], no widely used fitted model features the effect yet.

3.2.2 Aerosols. The second major part of an atmosphere are aerosols. As these particles are larger, scattering caused by them is described by Mie theory. An aerosol property database for these scatterers called OPAC (Optical Properties of Aerosols and Clouds) is available [Hess et al. 1998]: it contains several basic aerosol types that are typically present in an atmosphere, e.g. water-soluble (WASO), water-insoluble (INSO) and black carbon (SOOT). The optical properties of each of the types can be computed from the given particle size distributions and refractive indices and are available in libradtran [Emde et al. 2016].

The vertical profiles of the aerosol concentrations (Fig. 4, right) can be modeled for various environments according to OPAC, e.g. for continental (clean, average and polluted), urban, desert or arctic environments.

4 THE BRUTE FORCE SIMULATION

Goal of brute force atmospheric light transport is to generate authoritative images of complex, realistic simulated atmospheres. In this section, we give a technical overview of both the simulated atmosphere we use in the brute force renderings that serve as basis




Fig. 5. The effects of atmospheric ozone on sky dome appearance for a simplified atmosphere. **From left to right:** The fish-eye views from ground level show simulated atmospheres with no ozone, with a reduced ozone profile typical for ozone hole conditions (the green curve from Fig. 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.

of our model, as well as of `atmo_sim`, the path tracer used to obtain this reference dataset.

4.1 Viewing Distance as Model Parameter

In order to control the haziness of outdoor scenes, existing sky dome models feature a user-controllable parameter called *turbidity*. This is a relative measure of the fraction of additional scattering due to aerosols as opposed to molecules [Preetham et al. 1999], i.e., how much more the atmosphere scatters compared to an ideally clean molecular atmosphere:

$$T = \frac{t_m + t_h}{t_m}, \quad (1)$$

where t_m is the optical thickness of molecules only and t_h is the optical thickness of aerosols only. This value is typically reported at 550 nm and *measured towards the zenith*. As we will see in the next section, such vertical measurements are not a good correlate of haziness when using realistic scatterer profiles: they only correlate well when using exponential scatterer profiles.

As we are using realistic scatterer profiles with a large amount of aerosols close to the ground, we parameterise our model via *horizontal viewing distance at ground level (viewing distance for short)*. This is an intuitive parameter for end users, and we provide a model fit for viewing distances that range from 20 km to 130km.

4.2 Our Atmospheric Model

To achieve an appearance which is as realistic as possible, the composition of our atmospheric model is directly based on the U.S. Standard Atmospheres introduced in Sec. 3.2, with its profiles for three specific atmospheric configurations: *average*, *clean* and *polluted*. These standard profiles are based on real measurements, but simplified insofar as in the OPAC model these are modelled by two exponentials with a sharp transition at the inversion layer.

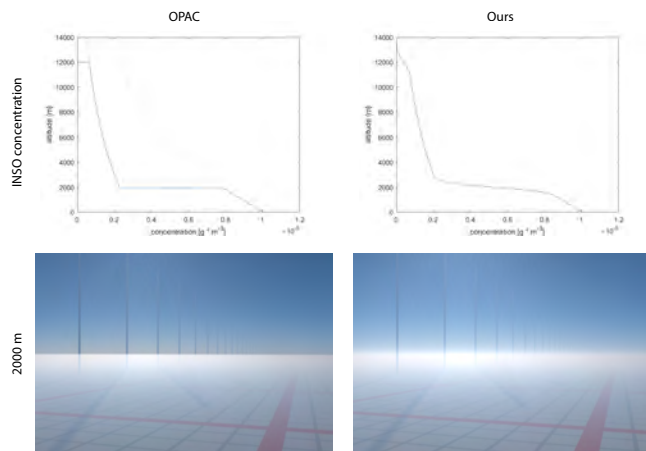


Fig. 6. A brute force comparison of the original OPAC atmosphere and our smoothed version, observer altitude 2000m. The sharp horizontal line seen in the left image is not the horizon, but the top of the denser scatterer region.

In order to obtain practically useful scatterer profiles, we first had to slightly modify the OPAC profiles by smoothing the inversion




Fig. 7. Comparison of our atmospheric model to real sky dome measurements by Kider et al. [2014]. **Left:** Photo from the Kider HDR dataset for 27-05-13, 09:30. **Middle:** Tone-mapped spectral render obtained with `atmo_sim` with continental clean aerosols, solar elevation 41.08°, ground albedo 0. **Right:** Spectral data difference plot using a similar colour scheme and scale as Kider et al., showing relative radiance error. The comparison directly uses the spectral data provided by Kider et al., and not the HDR images: the tone-mapped images shown here are included for illustrative purposes only. The left photo was post-processed, e.g., by removing black surroundings, flipping it to match the orientation of the spectral data, and colour shifting.

layer transition. For atmospheric research purposes, the discontinuities in the OPAC data likely do not matter: but as Fig. 6 shows, they are clearly visible for observer altitudes near the transition. Fig. 4 shows, in the rightmost plot, the smoothed vertical profiles we work with, and also the extrapolated profiles for viewing distances that were inserted to obtain smooth transitions for varying viewing distances. For all atmospheric configurations, we took WASO, INSO and SOOT particles into account. Their extinction cross sections and single scattering albedos were as provided by `libradtran`. The Rayleigh scattering cross section is based on Bodhaine et al. [1999] and only takes N_2 and O_2 into account. For O_3 particles, we use the cross sections from Gorshelev et al. [2014] at 223 K (Fig. 4, second plot).

In Sec. 3 of the supplementary PDF, we discuss the validation of our own simulations against results obtained for some of these atmospheric configurations with `libradtran`, and we also discuss the effect of slightly smoothing the OPAC reference data. Also, see Fig. 7 for a comparison against the measurements provided by Kider et al. [2014]. Our validations and comparisons show that our simulation provides results which are highly physically plausible, and that the curve smoothing eliminates undesirable artefacts.

4.2.1 Aerosol Phase Functions. For simplicity, the aerosol phase functions we use are based on a closed-form Henyey-Greenstein approximation [Henyey and Greenstein 1941], instead of full Mie scattering. As asymmetry parameters g for WASO and SOOT, we numerically fitted the best wavelength-dependent g to match the phase functions calculated by OPAC. For the strongly forward scattering INSO particles, our own fit resulted in a g value close to 1, which caused firefly artefacts in our brute force renderer, and caused a narrow and hard to fit high energy region around the solar disc to appear. For both these reasons, we resorted to using the INSO asymmetry parameters as defined by OPAC, which are slightly lower, and which slightly blur the circumsolar region. They are also responsible for the difference that can be seen in Fig. 7 in the immediate circumsolar region. For more details on the effects of using the INSO parameters we refer the reader to the `libradtran` validation in Sec. 3 of our supplementary PDF.

Table 2. 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.

	Sample values
Solar elevation (degrees)	-4.20, -4.00, -3.50, -3.00, -2.50, -2.00, -1.50, -1.00, -0.50, 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 (metres)	0.00, 1.87, 15.00, 50.62, 120.00, 234.38, 405.00, 643.12, 960.00, 1366.90, 1875.00, 2000.00, 2495.60, 3240.00, 4119.40, 5145.00, 6328.10, 7680.00, 9211.90, 10935.00, 12861.00, 15000.00
Ground albedo	0.00, 0.33, 0.66, 1.00
Viewing distance (km)	20, 27.6, 40.0, 59.4, 90.0, 131.8

4.3 The Dataset and Views used for Model Fitting

The dataset we use to fit the model consists of renderings which systematically cover the desired parameter space. Inputs are solar elevation at ground level, observer altitude, ground albedo, and viewing distance. Values for the parameters can be found in Table 2. We compute the spectral model for the same 11 wavelengths as [Hošek and Wilkie 2012]. We experimented with changing this choice of spectral channels, but found that fewer started to cause noticeable colour shifts, while more yielded no tangible benefits.

Our goal is a model for computer graphics (i.e., for the human visual system, with its fairly broadband sensors), and for the training of machine vision applications (which typically also involve sensors with fairly broad spectral characteristics), so this fairly low number of spectral bands is not a concern. If nanometer-level spectral accuracy were desired, one would need to resort to direct computations via `libradtran`, or a similar specialised tool: pre-computing a sky-dome model to such a fine-grained spectral resolution would require prohibitive amounts of storage space.

4.3.1 Side-facing Fish-Eye Views. One significant difference to previous work is that our model now contains the ability to describe in-scattered radiance for varying observer altitudes. Therefore, except at ground level, this means we now have to fit a fully spherical function, and not a hemi-spherical one. In order to avoid discontinuities at the horizon, we use a single side-facing fish-eye view of the sky as input to the fitting algorithm, and not two fish-eye views that face up and down. For the idealised planet defined in the brute force tracer, the in-scattering data is symmetrical left and right of the solar position, so we only have to generate a single side-facing view.

Fig. 8 shows some side-facing views for one particular viewing distance and solar elevation: the increasing influence of in-scattered light in the lower half of the hemisphere can be seen, as well as the curvature of the Earth that appears for higher observer altitudes.




Fig. 8. Side-facing fish-eye views of in-scattered radiance similar to those used in fitting the model. Solar elevation 8° , ground albedo 0.3, sun at the right. Observer altitudes (from left to right) 0, 3.5, 7, 15 and 75 km. The latter is far above the altitudes used to fit the model, and shows that the used simulator is also capable of rendering such “views from space”. **Top row:** tone mapped images. The ground is black in the left-most image, as only in-scattered light is stored – ground reflections are intentionally omitted. **Bottom row:** scaled overlay polarisation visualisation of the degree of polarisation at 550 nm, provided by `polvis` (see Sec. 4.4.1).

4.3.2 Solar Elevation. The uneven sampling of solar elevation near the horizon, near the ground and near the inversion layer (based on the observer altitude) was determined iteratively through analysing test renderings, and examining how fast sky dome features change for varying input parameters in each region of the parameter space.

4.4 The Atmospheric Path Tracer

Generating reference datasets to fit our model to requires an unbiased brute force renderer which is capable of computing large amounts of spectral images for our chosen atmospheric configurations. These input images for our fitting need to contain polarisation information, and the used rendering algorithm should be flexible enough that we can exclude directly viewed ground hits – we only want to fit in-scattered radiance that originates from atmospheric events. In addition to these requirements, we wanted the brute force renderer to be able to load scene geometry, to run direct brute force computations for non-trivial scenes.

With regard to accuracy of radiative transfer computations in the atmosphere, the `libradtran` software package [Emde et al. 2016] is the yardstick to measure against. However, while `libradtran` is an excellent tool for reference computations, it would not have been well suited to compute the large sets of images we need for fitting our model. Part of the efficiency of `libradtran` comes from the fact that it always runs only for a single wavelength, and a single query direction: this makes limited bi-directional tracing (and therefore faster convergence) considerably easier. However, we need entire reference images that contain several fairly broad spectral bands: and repeatedly running `libradtran` for multiple query directions and wavelengths is not very efficient. Also, we are not aware of an option that allows one to only compute in-scattered energy from atmospheric events, and omit specific ground hits: so we would likely have had to significantly modify their Monte Carlo estimator.

4.4.1 `atmo_sim`. Instead, we wrote our own brute force simulator `atmo_sim`, based on the ART framework [Wilkie 2018]. ART provides infrastructure for storing, analysing and manipulating spectral images that contain polarisation information, and it includes the

command line tool `polvis` which we used for the polarisation visualisations in Figs. 8, 10 and 15. `atmo_sim` is an uni-directional path tracer which uses multiple importance sampling [Veach 1997], and it is optimised to deal with scattering events in an atmosphere around an idealised Lambertian planet. Via Null Scattering [Miller et al. 2019] in combination with Hero wavelength sampling [Wilkie et al. 2014], it achieves a rendering performance which is sufficient for the generation of reference datasets, even for below sunset scenarios: and we verified the results obtained with it against `libradtran` (see Sec. 3 of the supplementary PDF for a discussion of this).

5 CREATION OF THE ANALYTICAL MODEL

In the following sections we discuss the components of the model: in-scattered radiance, the in-scattered polarisation component, and atmospheric transmittance.

5.1 In-scattered Radiance

As the inner workings of the image-based decomposition we introduce for representing in-scattered radiance are somewhat complex, we give a detailed explanation of it in Sec. 1 of the supplementary PDF, and only present the main characteristics of the approach here.

Because of the large number of appearance features that are present on a fully spherical sky dome, especially if post-sunset scenarios are included (see Fig. 9 for examples of these), we had to develop an entirely new fitting approach that is distinct from previous techniques. In our model, we obtain the radiance pattern of the sky as a sum of outer products of single variable functions. The functions themselves are free-form, tabulated and were obtained by Canonical Polyadic Decomposition (CPD) [Kolda and Bader 2009], a process very similar to SVD low rank approximation. This approach can be thought of as a specialised compression scheme, however it is also essentially a decomposition of the radiance pattern into an optimal orthogonal set of “features”; the methods that we work with all rely on tensor and matrix decompositions.

The performance of CDP critically depends on using a suitable input parameterisation that allows the separation to take place cleanly: and in this regard, we found a suitable scheme, described in Sec. 1.2 of the supplementary PDF, that is based on the solar angle plus the shadow and zenith angles, and which makes the gradient of the solar glow and the shadow/horizon lines aligned with both axes of the fish-eye input images that are re-projected to this space. There, they are expressed as an outer product of two vectors using the CPD decomposition, and we store the two vectors for later retrieval and reconstruction.

Subtle changes to sky dome appearance both with changing observer altitude and with the sun going beneath the horizon actually require further refinements to this basic idea, like the high altitude correction and the image pre-emphasis to improve horizon interpolation and decomposition of post sunset images which are described in Sec. 1.3 and Sec. 1.5 of the supplementary PDF. It has to be noted that while the basic idea of using CDP to handle a dataset like ours is conceptually simple, the approach would not have yielded a useful result without these further refinements that are specific to the problem that we wish to model.

5.2 Polarisation Component

Building an unoptimised analytical model for sky dome polarisation would require to fit functions to the three additional Stokes components present in each polarised wavelength sample. However, sky dome polarisation is almost entirely linear, which reduces the problem to fitting the second and third components: together, these encode the entire linear polarisation component in a Stokes vector. And for sky domes, one can take this even further: as illustrated in Fig. 10, one can, as a post-process, 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-parameterisation, one can get a good approximation for sky dome polarisation patterns by only fitting the second Stokes component. The polarisation 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 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 was actually permissible. We ran the entire brute force simulation with polarisation enabled, so we are able to actually assess the impact of this simplification. While it does introduce an error insofar as the polarisation effects of higher order scattering are discarded, Fig. 10 shows that after re-parameterisation, the third channel only contains a weak signal that has no perceptible effect, even in scenes that exhibit visual differences due to sky-dome polarisation. 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 polarisation pattern is the same as for the main radiance model: however, as the pattern is much simpler, we can reduce the decomposition rank to $n = 5$, and also reduce the sizes of the individual function tables.

5.3 Atmospheric Transmittance

Rendering of outdoor scenes requires the computation of atmospheric transmission of light between two points. As we assume an atmospheric configuration that is rotationally symmetric with respect to the surface normal, the transmission for a given atmosphere configuration is a 2D distribution that varies with wavelength and observer altitude.

To create a fitted model of this function, we first precomputed atmospheric transmission in a coordinate system aligned with the curvature of the planet for the same set of altitudes and wavelengths as the in-scattered radiance model. We then used Singular Value Decomposition (SVD) to produce a low rank approximation of the transmittance. A separate fit is computed for each altitude, where the SVD is computed over all atmospheric configurations and wavelengths at that altitude. Transmittance between two points can be computed by projecting one of the points into the coordinate system used for transmittance, reconstructing from the low rank approximation, then interpolating this value between altitudes and wavelengths. Please see Sec. 2 of the supplementary PDF for more details. Fig. 11 shows the transmittance component of a scene computed with Monte Carlo and our fit. This shows our method is able




Fig. 9. Luminance patterns for below sunset conditions with continental average aerosols. **From left to right:** Solar elevations 0° to -6° . **Top row:** Images individually tone mapped for similar overall brightness. Note how the shadow of the Earth rises opposite the location where the sun sets below the horizon, and that the shadow is less and less distinct as night sets in. **Bottom row:** The same images tone mapped to equal the brightness of solar elevation 0° .

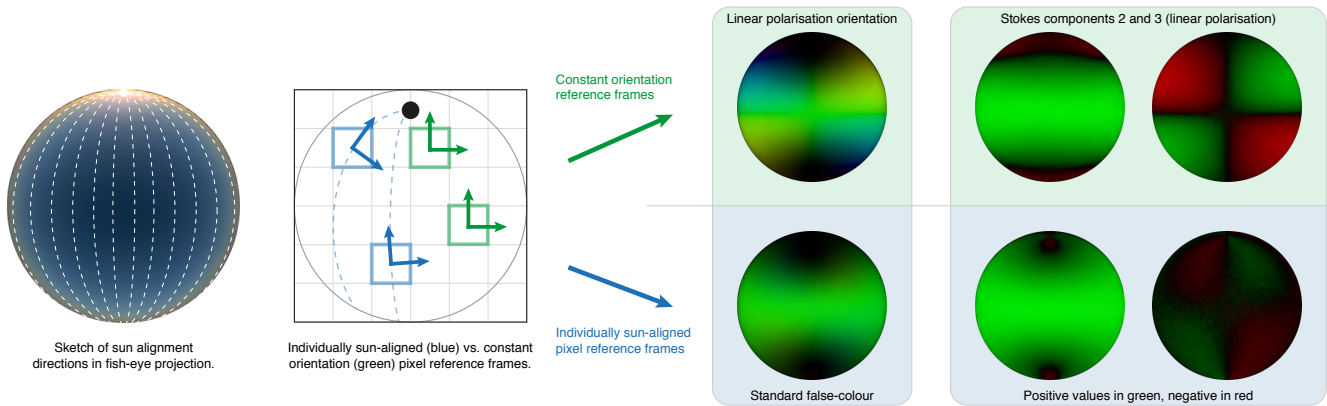


Fig. 10. Illustration of the sun-aligned polarisation reference frames which allow us to only work with the second Stokes component of polarised sky dome images. Standard reference frames are shown in green, individually aligned ones in blue. The reference frames are right-hand coordinate systems, as the light is coming from the image plane towards the observer. The colour scheme used in the “Standard false-colour” images is a linear polarisation orientation plot provided by polvis (see Sec. 4.4.1). Note that even with the special alignment, the third Stokes component is not zero, but just very weak: this component is caused by multiple polarising scattering events, which are rare. 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. Note that in this figure we demonstrate the concept with upward-facing fish-eye views, as these give a view of the complete polarisation 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.

to efficiently model atmospheric transmittance for a wide range of distances.

5.4 Finite In-scattered Radiance Calculation

Given that our model allows one to calculate radiance for infinite paths and transmittance for finite paths, in-scattered radiance between observer x_1 , viewed point x_2 and view direction $\vec{v} = x_2 - x_1$ can be calculated as

$$L_{in}(x_1, x_2) = L_{\infty}(x_1, \vec{v}) - \tau(x_1, x_2) * L_{\infty}(x_2, \vec{v}) \quad (2)$$

where $L_{in}(x_1, x_2)$ denotes in-scattered radiance arriving at point x_1 along a given finite path from point x_2 , $\tau(x_1, x_2)$ denotes the transmission between the two points, and $L_{\infty}(x_1, \vec{v})$ denotes in-scattered radiance along a given infinite path coming from direction

\vec{v} . This formula also works for all components of polarised radiance in a Stokes Vector.

For practical use, it has to be noted that equation 2 only holds when the two L_{∞} values used in it are both highly numerically accurate. However, in a fitted model such as ours, for any two distinct query locations, there can always be small radiance discrepancies compared to the ground truth. In rendered images, these deviations will manifest themselves as horizontal stripe artefacts for finite viewing distances, usually for observer altitudes near the ground, and at viewing angles near the horizon.

To solve these stability issues, one has to consider the effect our compression has on sky-dome radiance patterns near the horizon. As shown in the supplementary PDF, the most visible change compared to the original brute force images is that our model slightly blurs the




Fig. 11. Visualisations of the transmission component for an observer altitude of 8 km, for the test scene shown in Fig. 14. The left image is a reference rendering obtained by brute force path tracing, and the right is the result obtained by using our fitted SVD approximation.

horizon region. What is not immediately obvious is that it does so within a small and not entirely predictable blur range. For directly observed radiance, these small variations in blur are imperceptible: but for the above mentioned viewing geometries, equation 2 still suffers numerical stability issues because of it.

As there is no way to get rid of the existing variable blur near the horizon region, a workable alternative is to *always actively bring all such lookups to a consistent minimum level of blurriness*. Equation 2 becomes stable if the data does not exhibit small random variations due to compression: and this can easily be achieved by not taking a single sample in the exact path direction, but the average of a few directions slightly above and below $x_2 - x_1$ instead.

6 DISCUSSION AND RESULTS

When integrating our new model in actual path tracing software for testing purposes, we faced two main obstacles. The first was that evaluation of a pre-computed sky dome model which features a full spherical radiance fit that changes with observer altitude requires more than just a drop-in replacement of whatever sky dome model is already present in a given path tracer: making full use of the altitude-dependent capabilities of the model requires modifications to the light source sampling code of the target system. To our knowledge, all existing renderers assume hemispherical models that do not change with altitude when dealing with analytical sky dome radiance. The second obstacle was that in order to use our model, which is fully spectral, we needed a target rendering system which is also spectral, and optimally even polarisation capable.

The spectral requirements narrowed the field down considerably. Mitsuba 2 [Nimier-David et al. 2019] is fully spectral and is polarisation capable, but was only comparatively recently released to the general public. The only other rendering system with such capabilities is ART [Wilkie 2018], which is however quite limited with regard to the scene types it can actually render. We first experimented with including the model in Mitsuba 1, which worked well. But as our model will now also be included in the next release of the Corona renderer [Chaos Czech a.s. 2021], we used that system to generate Fig. 1 instead. And to give an indication of polarisation performance in an actual renderer, we also included it in ART: we opted to use that instead of Mitsuba 2, as `atmo_sim` is based on ART, so we had prior experience with it. We donated code and data to the

ART project, so a full implementation of our model will be available as Open Source in the next release of that system.

6.1 Size of Our Model

As described in Sec. 4.2, we fit our model for 6 viewing distances, each of which yields a dataset of 283MB for just radiance and transmission data, or 410 MB if polarisation data is included. The entire polarisation-enabled model for ground level viewing ranges from 20 km to 130 km is 2.4GB in size, or 1.66 GB for just radiance and transmission.

6.2 Performance Considerations

With its added capabilities such as fully spherical radiance patterns, transmission and in-scattering for finite distances, the rendering performance of our model is not easily comparable to previous models like Preetham or Hosek. When the in-scattering computations for finite distances are used, a renderer using the model of course runs slower than with e.g. pure Hosek sky dome look-ups, and no volumetric computations. But if our model is used to only provide sky dome radiance, it yields results that are qualitatively similar to, but still more realistic than, the Hosek model: see Fig. 3 for an example of this. In our experiments we used both models in exactly the same renderer: and we found that they run at practically the same speed. For instance, for two typical renderings shown in this work (Fig. 15, and the Villa Rotonda haze render found in the supplementary materials), render times were e.g. 228 vs. 208 seconds (Hosek was slightly faster), and 124 vs. 126 seconds (slight




Fig. 12. An example of how the results of the brute force renderer (left column) and the fitting (right column) match, for in-scattered radiance (top row) and for the sun-aligned (cf. Fig. 10) first polarisation component (bottom row). Solar elevation 18.64°, observer altitude 2495m.

advantage to our model), respectively. And that was with the more complex, altitude-dependent light source evaluation code running for our model: so the actual raw model queries are definitely faster than Hosek.

Our current implementation loads the entire dataset into memory upon starting the renderer. Therefore it requires at least 600 MB RAM to run. However, for scenes with a limited range of altitudes only a subset of the full dataset will be used which allows significant memory savings, if memory consumption were an issue.

6.3 Aerial Perspective

Figs. 1, 14 show results for the combination of sky dome radiance with attenuation and in-scattered light for finite viewing distances, as discussed in Sec. 5.4. In such a setting, the fitted model is significantly faster than the brute force path tracing of the atmospheric model it replaces, which required tens of thousands of samples per pixel for convergence. In contrast, the spectral path tracer we integrated the fitted model into, requires at most 256-512 SPP to converge to virtually noise-free estimates of directly viewed diffuse surface. Which is a fairly typical value for such path tracers and complex environment lighting, so the sky dome model itself is not a constraining factor for image convergence.

For the scene shown at a resolution of 1500×1000 in Fig. 14, the statistics are as follows: we used up to 85k spp and 5000 core hours (= 10 days on 24 core CPU) for the reference renders, with the post-sunset case being the most difficult one. By comparison, the renders that used the fitted model used 100 spp and 2 core hours each (= 5 mins on 24 core CPU).

6.4 Polarisation

As ART has limited modelling capabilities and cannot use mainstream scene description formats, and also for an equal comparison with [Wilkie et al. 2004], we obtained the original scene file for one of their test scenes from the authors of [Wilkie et al. 2004], resulting in Fig. 15. Inter-reflections between skyscrapers, or glass facades and water bodies in outdoor settings can be considerably mis-predicted by a non-polarising renderer. As our model includes in-scattered radiance for finite viewing directions, it is capable of effects such as reproducing the scenario that a part of the in-scattered light in a daytime outdoor scene can be removed with a polarisation filter: a part of the haze can be “lifted” that way. The supplementary materials include a pair of EXR images that show this effect.

6.5 Comparison to Other Models

As the use of exponential scatterer distributions is a common feature of many other models, we show further comparisons of our model against an exponential distribution and a reference render computed using the brute force path tracer in Fig. 13. These results show similarities between the models at lower altitudes, however these models diverge as altitude increases due to the realistic non-exponential scatterer distribution used by our work. Furthermore, we directly compare to the sky models from Hillaire [2020] and Bruneton [2016] in Sec. 7 of the supplementary PDF.

7 CONCLUSION AND FUTURE WORK

We have introduced a comprehensive and realistic pre-computed sky dome model for rendering outdoor environments under clear skies. It improves the state of the art in several ways:

- (1) The model is based on OPAC standard atmospheric constituent data, so the molecular and aerosol distributions represent realistic viewing conditions.
- (2) The novel in-scattered radiance model provides a full spherical fitting, and is available for observer altitudes up to 15 km altitude. This allows viewpoints above ground level, and downward looking renders. We also show how to evaluate our new in-scattering model for finite viewing distances.
- (3) The model extends to solar elevations past sunset, and provides post-sunset features such as proper twilight blue (due to ozone absorption in the high atmosphere), and the shadow of the Earth rising opposite the setting sun.
- (4) We provide a matching transmission function that is view-direction dependent, and also changes with observer altitude.
- (5) We also provide a fitted function for the linear polarisation patterns found in clear skies.

As these components are all derived from the same ground truth dataset, they are consistent, and can be seamlessly used together. To our knowledge, features 1 to 3 are true novelties, and are not available in any existing pre-computed models of sky dome radiance. Feature 4 is only available for [Preetham et al. 1999], but not in any later models. Feature 5 was attempted by [Wilkie et al. 2004] and [Wang et al. 2016], but our integrated model is the first to provide polarisation information which matches all other components. The only limitations of our new model are the size of the dataset, and the fact that hazier atmospheric configurations are an extrapolation of OPAC data.

In the future, we plan to improve our model in the following areas:

- (1) Most importantly, the model can easily be made to use genuine scattering profile data for high haziness situations – iff such data is found in atmospheric science literature. A key issue here seems to be purely meteorological, in the sense that there are potentially a large number of different fog and haze configurations that can cause low ground level visibility: and merely having a single “view distance” parameter is not enough to differentiate between these.
- (2) Special provisions can be made to accommodate the sharp circumsolar features that appear if “real” INSO lobes were to be used, instead of the smoothed OPAC ones.
- (3) More subtle polarisation patterns could be represented if full Mie scattering were used in the brute force simulator: however, this would substantially increase the size of the polarisation fitting, as the optimisation discussed in Sec. 5.2 could no longer be used.

Finally, the current cut-off altitude of 15 km for the observer altitude is arbitrary, and could easily be extended higher, if this was needed.




Fig. 13. A comparison of our OPAC-based atmosphere scatterer profile versus a purely exponential one. The key difference is that no clear sense of observer altitude can manifest itself with an exponential fall-off: in a real clear atmosphere, there is a distinct hazier layer in the first 1-3km from the ground, with significantly clearer air above it. A purely exponential model is therefore unable to provide realistic views from mountaintops or aircraft, where this feature plays an important role in overall scene appearance. The supplementary PDF contains further comparisons between exponentials and OPAC-derived scatterer profiles.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their insightful comments, which were a considerable help on the somewhat protracted journey this paper had towards publication. We acknowledge funding by the Czech Science Foundation grant 19-07626S, the EU H2020-MSCA-ITN-2020 grant number 956585 (PRIME), and the Charles University grant number SVV-260588. This work was supported through the Intel Graphics and Visualization Institute at Charles University. This work was supported by The Czech Ministry of Education, Youth and Sports from the Large Infrastructures for Research, Experimental Development and Innovations project "IT4Innovations National Supercomputing Center – LM2015070". The authors also acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing HPC resources that have contributed to the research results reported within this paper. Chaos Czech a.s. provided the project with the city scene used in Fig. 1.

REFERENCES

- G. Anderson, Shepard Clough, F. Kneizys, J. Chetwynd, and Eric Shettle. 1986. AFGL Atmospheric Constituent Profiles (0.120km). (05 1986), 46.
- Barry A. Bodhaine, Norman B. Wood, Ellsworth G. Dutton, and James R. Slusser. 1999. On Rayleigh Optical Depth Calculations. *Journal of Atmospheric and Oceanic Technology* 16, 11 (1999), 1854–1861. [https://doi.org/10.1175/1520-0426\(1999\)016<1854:ORODC>2.0.CO;2](https://doi.org/10.1175/1520-0426(1999)016<1854:ORODC>2.0.CO;2) arXiv:[https://doi.org/10.1175/1520-0426\(1999\)016<1854:ORODC>2.0.CO;2](https://doi.org/10.1175/1520-0426(1999)016<1854:ORODC>2.0.CO;2)
- Eric Bruneton. 2016. A qualitative and quantitative evaluation of 8 clear sky models. *IEEE transactions on visualization and computer graphics* 23, 12 (2016), 2641–2655.
- Eric Bruneton and Fabrice Neyret. 2008. Precomputed atmospheric scattering. In *Computer graphics forum*, Vol. 27. Wiley Online Library, 1079–1086.
- Chaos Czech a.s. 2021. Corona Renderer. <https://corona-renderer.com/>.
- C. Emde, R. Buras-Schnell, A. Kylling, B. Mayer, J. Gasteiger, U. Hamann, J. Kylling, B. Richter, C. Pause, T. Dowling, and L. Bugliaro. 2016. The libRadtran software package for radiative transfer calculations (version 2.0.1). *Geoscientific Model Development* 9, 5 (2016), 1647–1672. <https://doi.org/10.5194/gmd-9-1647-2016>
- V. Gorshelev, A. Serdyuchenko, M. Weber, W. Chehade, and J. P. Burrows. 2014. 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 (2014), 609–624. <https://doi.org/10.5194/amt-7-609-2014>
- David Guimera, Diego Gutierrez, and Adrián Jarabo. 2018. A Physically-Based Spatio-Temporal Sky Model. In *Spanish Computer Graphics Conference (CEIG)*, Ignacio García-Fernández and Carlos Ureña (Eds.). The Eurographics Association. <https://doi.org/10.2312/ceig.20181150>
- Jörg Haber, Marcus Magnor, and Hans-Peter Seidel. 2005. Physically-based simulation of twilight phenomena. *ACM Transactions on Graphics* 24 (October 2005), 1353–1373. Issue 4. <https://doi.org/10.1145/1095878.1095884>
- Miles Hansard. 2019. Fast Synthesis of Atmospheric Image Effects. In *European Conference on Visual Media Production* (London, United Kingdom) (CVMP '19). Association for Computing Machinery, New York, NY, USA, Article 13, 10 pages. <https://doi.org/10.1145/3359998.3369406>
- L. G. Henyey and J. L. Greenstein. 1941. Diffuse radiation in the Galaxy. *Astrophysical Journal* 93 (Jan 1941), 70–83. <https://doi.org/10.1086/144246>
- M. Hess, P. Koepke, and I. Schult. 1998. Optical Properties of Aerosols and Clouds: The Software Package OPAC. *Bulletin of the American Meteorological Society* 79, 5 (1998), 831–844. [https://doi.org/10.1175/1520-0477\(1998\)079<0831:OPOAAC>2.0.CO;2](https://doi.org/10.1175/1520-0477(1998)079<0831:OPOAAC>2.0.CO;2) arXiv:[https://doi.org/10.1175/1520-0477\(1998\)079<0831:OPOAAC>2.0.CO;2](https://doi.org/10.1175/1520-0477(1998)079<0831:OPOAAC>2.0.CO;2)
- Sébastien Hillaire. 2020. A Scalable and Production Ready Sky and Atmosphere Rendering Technique. *Comput. Graph. Forum* 39, 4 (2020), 13–22.
- Lukas Hošek and Alexander Wilkie. 2012. An analytic model for full spectral sky-dome radiance. *ACM Trans. Graph.* 31, 4 (2012), 95. <http://doi.acm.org/10.1145/2185520.2185591>
- Lukas Hošek and Alexander Wilkie. 2013. Adding a solar radiance function to the Hošek-Wilkie skylight model. *IEEE Computer Graphics and Applications* 33, 3 (2013), 44–52. <http://dx.doi.org/10.1109/MCG.2013.18>
- E. O. Hulburt. 1953. Explanation of the brightness and color of the sky, particularly the twilight sky. *J. Opt. Soc. Am.* 43 (1953), 113–118.
- AK Kaifel, M Felder, C DeClercq, and J-C Lambert. 2012. New dynamic NNORSY ozone profile climatology. *Atmospheric Measurement Techniques Discussions* 5, 1 (2012), 775–812.
- Joseph T. Kider, Jr., Daniel Knowlton, Jeremy Newlin, Yining Karl Li, and Donald P. Greenberg. 2014. A Framework for the Experimental Comparison of Solar and Skydome Illumination. *ACM Trans. Graph.* 33, 6, Article 180 (Nov. 2014), 12 pages. <https://doi.org/10.1145/2661229.2661299>

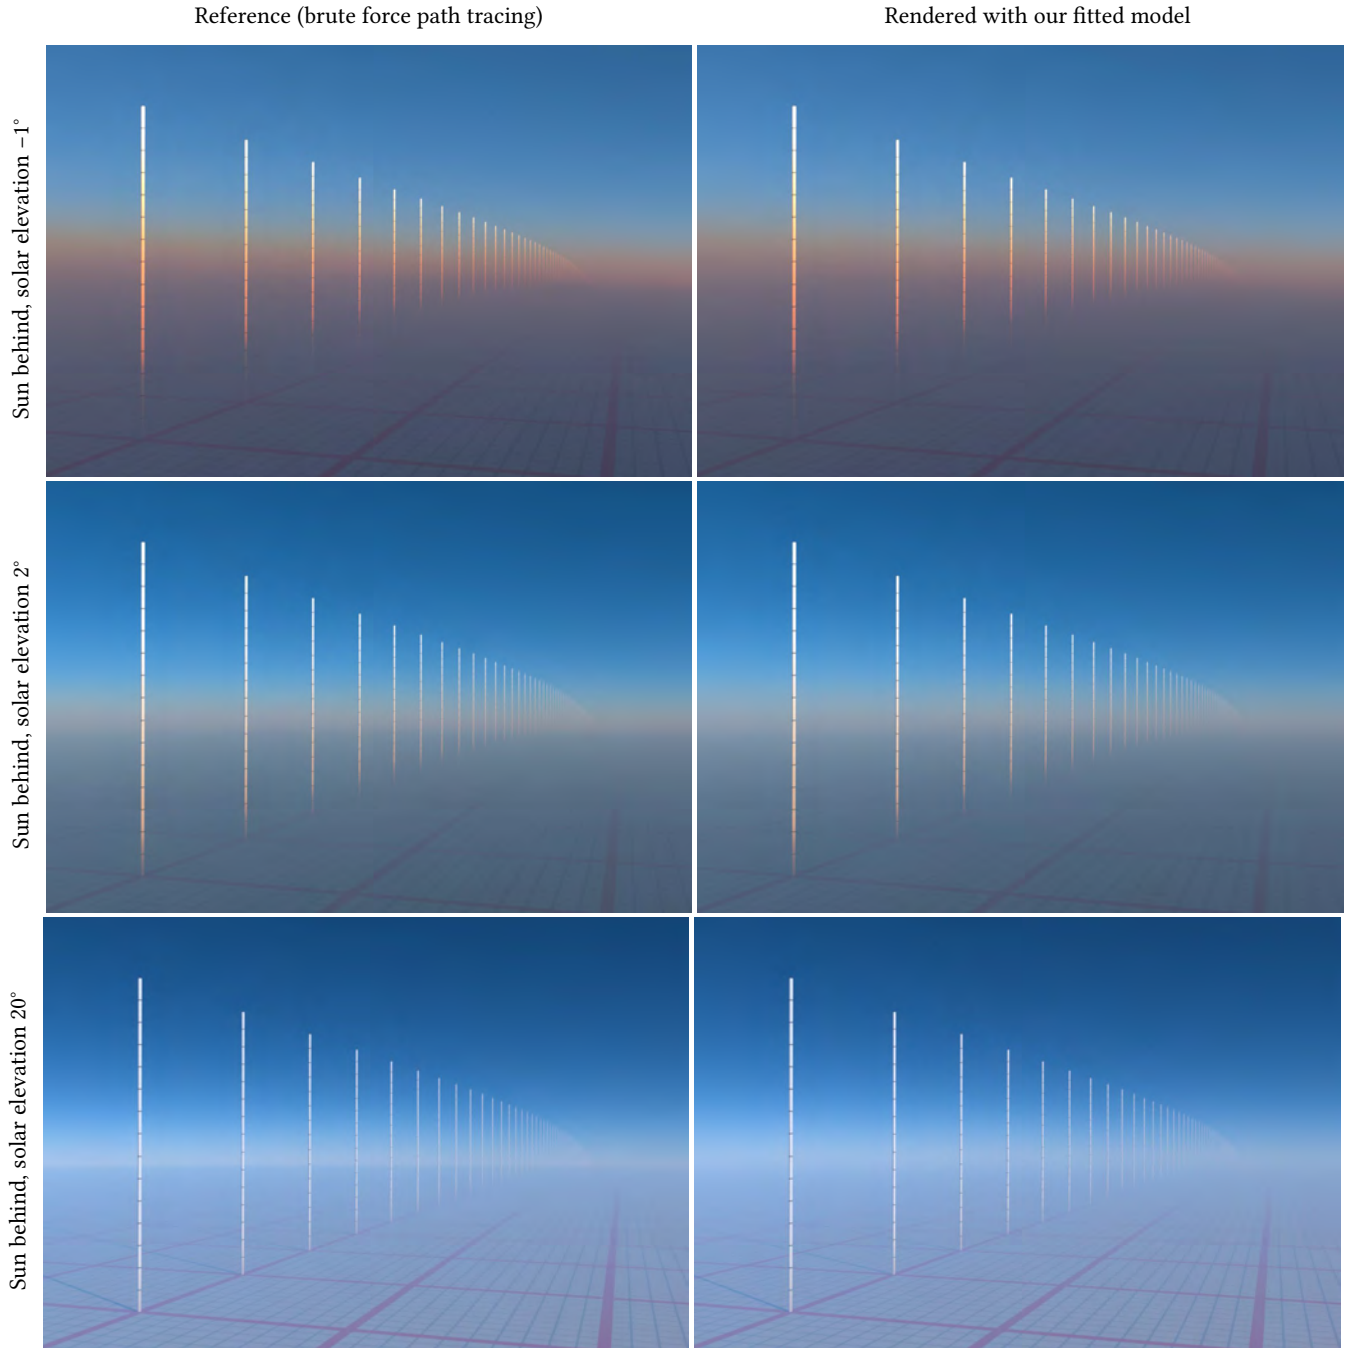


Fig. 14. Comparison of reference brute force path tracing results from `atmo_sim` (left) to renderings using our pre-computed model (right) at solar elevations -1° , 2° and 20° (from top to bottom). The test scene contains a planet with a 10×10 km red grid texture, with 15 km high towers with stripes every 1 km (cf. Figure 16 in the supplementary PDF for scale). The camera is at 8 km altitude, and the sun is behind the camera. Note that the reference dataset did not contain solar elevations 2° and 20° , which shows that our model is capable of interpolations that match the references. The only apparent difference can be seen at the horizon, where our model is more blurred, which can also be seen in Fig. 12 and is mainly caused by the limited resolution of the fit. Note that this blurring mainly affects extreme viewing distances over perfectly flat terrain, and should not lead to artefacts in normal scenes. Finally, while the path tracer uses the true spatially-varying surface albedo, our model is parametrised by a single average value leading to a subtle colour shift on the ground in the last image.




Fig. 15. The specular architecture scene from Wilkie et al. [2004], re-created using ART [2018]. Qualitatively similar behaviour as seen in the 2004 publication can be observed. Image rendered with a polarising skylight (top left), and a non-polarising version of our model (top right). The difference image (bottom left) shows that skylight polarisation not only affects the reflection of the sky in the building facades, but also the inter-reflections of the buildings. The image on the bottom right shows the degree of polarisation as scaled overlay for 550 nm, provided by polvis (see Sec. 4.4.1).

- T. Kolda and B. Bader. 2009. Tensor Decompositions and Applications. *SIAM Rev.* 51, 3 (2009), 455–500.
- Peter Kutz. 2013. Sky Renderer project blog. <http://skyrenderer.blogspot.com>. Accessed: 2015-12-31.
- Raymond L. Lee, Wolfgang Meyer, and Goetz Hoeppe. 2011. Atmospheric ozone and colors of the Antarctic twilight sky. *Applied Optics* 50, 28 (2011), 162–171.
- Christian Mätzler. 2002. MATLAB Functions for Mie Scattering and Absorption. *Research report 2002-8, Institut für Angewandte Physik, Universität Bern, Switzerland* (2002). <http://www.iap.unibe.ch/publications/download/201/en/>
- Bailey Miller, Iliyan Georgiev, and Wojciech Jarosz. 2019. A null-scattering path integral formulation of light transport. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–13.
- Merlin Nimier-David, Delio Vicini, Tizian Zeltner, and Wenzel Jakob. 2019. Mitsuba 2: A Retargetable Forward and Inverse Renderer. *ACM Trans. Graph.* 38, 6, Article 203 (Nov. 2019), 17 pages. <https://doi.org/10.1145/3355089.3356498>
- Tomoyuki Nishita, Yoshinori Dobashi, and Eihachiro Nakamae. 1996. Display of clouds taking into account multiple anisotropic scattering and sky light. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*. 379–386.
- Tomoyuki Nishita, Takao Sirai, Katsumi Tadamura, and Eihachiro Nakamae. 1993. Display of the earth taking into account atmospheric scattering. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques (Anaheim, CA) (SIGGRAPH '93)*. ACM, New York, NY, USA, 175–182. <https://doi.org/10.1145/166117.166140>
- S O'Neal. 2005. Accurate Atmospheric Scattering. GPU Gems 2.
- R. Perez, R. Seals, and J. Michalsky. 1993. All-weather model for sky luminance distribution—Preliminary configuration and validation. *Solar Energy* 50, 3 (1993), 235 – 245. [https://doi.org/DOI:10.1016/0038-092X\(93\)90017-I](https://doi.org/DOI:10.1016/0038-092X(93)90017-I)

- Matt Pharr and Greg Humphreys. 2010. *Physically Based Rendering, Second Edition: From Theory To Implementation* (2nd ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- A. J. Preetham, Peter Shirley, and Brian Smits. 1999. A practical analytic model for daylight. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques (SIGGRAPH '99)*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 91–100. <https://doi.org/10.1145/311535.311545>
- Chandrasekhar Subrahmanyam. 1960. Radiative transfer.
- HC van de Hulst. 1957. *Light scattering by small particles*. Dover Publications.
- Eric Veach. 1997. *Robust Monte Carlo methods for light transport simulation*. Ph.D. Dissertation. Stanford University, Stanford, CA, USA.
- Xin Wang, Jun Gao, Zhiguo Fan, and Nicholas W Roberts. 2016. An analytical model for the celestial distribution of polarized light, accounting for polarization singularities, wavelength and atmospheric turbidity. *Journal of Optics* 18, 6 (2016), 065601. <http://stacks.iop.org/2040-8986/18/i=6/a=065601>
- Alexander Wilkie. 2018. The Advanced Rendering Toolkit. <http://cgg.mff.cuni.cz/ART>.
- Alexander Wilkie, Sehera Nawaz, Marc Droske, Andrea Weidlich, and Johannes Hanika. 2014. Hero wavelength spectral sampling. In *Computer Graphics Forum*, Vol. 33. Wiley Online Library, 123–131.
- Alexander Wilkie, Christiane Ulbricht, Robert F. Tobler, Georg Zotti, and Werner Purgathofer. 2004. An Analytical Model for Skylight Polarization. In *Rendering Techniques*, Alexander Keller and Henrik Wann Jensen (Eds.). Eurographics Association, 387–398.