

A Wide Spectral Range Sky Radiance Model

P. Vévoda^{1,2}, T. Bashford-Rogers³, M. Kolářová¹ and A. Wilkie¹

¹Charles University, Czech Republic

²Chaos Czech, Czech Republic

³University of the West of England, United Kingdom




Figure 1: Our spectral sky model provides radiance for a wide range of wavelengths. It extends the Prague Sky Model [WVBR*21] and adds another 44 wavelength channels in the near-UV and near-IR range on top of its original 11 channels in the visible range. Altogether, our model covers wavelengths from 280 nm to 2480 nm using 55 regularly spaced bins. The displayed monochrome images correspond to 280, 1200, 1600, 2000 and 2480 nm bins; the colour image is a composition of 11 bins from 320 – 760 nm range. All images show the same model parameters: visibility 20.0 km, solar elevation 3 °, and ground albedo 0.3.

Abstract

Pre-computed models of sky radiance are a tool to rapidly determine incident solar irradiance in applications as diverse as movie VFX, lighting simulation for architecture, experimental biology, and flight simulators. Several such models exist, but most provide data only for the visible range and, in some cases, for the near-UV. But for accurate simulations of photovoltaic plant yield and the thermal properties of buildings, a pre-computed reference sky model which covers the entire spectral range of terrestrial solar irradiance is needed: and this range is considerably larger than what extant models provide. We deliver this, and for a ground-based observer provide the three components of sky dome radiance, atmospheric transmittance, and polarisation. We also discuss the additional aspects that need to be taken into consideration when including the near-infrared in such a model. Additionally, we provide a simple standalone C++ implementation as well as an implementation with a GUI.

CCS Concepts

• *Computing methodologies* → *Rendering*; • *Applied computing* → *Earth and atmospheric sciences*;

1. Introduction

Accurate simulation of light transport in realistic atmospheres has been possible for a considerable time now. However, in the past, and at least within computer graphics, the computational cost of such simulations was prohibitive for routine use: and even now they are very demanding, so that usually one still resorts to approximations of some sort.

Outright simulations of atmospheric light transport being infeasible for production work led to numerous pre-computed approximations of the radiance patterns observed on clear skies, and direct solar irradiance, being developed. These were initially quite simple, like the hemispherical luminance-only model developed by Perez [PSM93], but later progressed to fully spherical spectral models with a polarisation component, like the Prague Sky Model [WVBR*21]. Such models are now routinely used in ap-

plication areas such as lighting design and even for rough estimates of photovoltaic plant yield.

However, all extant pre-computed models were mainly designed with the visible range, and a human observer, in mind. So, specifically for the application cases of photovoltaic plant yield simulation and thermal analysis of buildings, they lack a significant part of the solar irradiance spectrum: specifically, the near-infrared region up to around 2500 nm is missing. With the importance of renewable energy sources rapidly increasing, and the thermal performance of buildings also becoming more and more important, accurate design tools for these applications are urgently needed: even though these application areas developed their own ad hoc prediction toolchains over the years (some of which we discuss in Section 2), the trend is now to move towards the kind of Monte Carlo (MC) rendering technology that was originally developed for “normal” computer graphics. The reason being that only MC rendering can handle predictions for arbitrarily complex input geometries, non-trivial surface materials, transparency and translucency (for e.g. complex photovoltaic module coverings), and similar advanced appearance features. And as computer graphics technology is the best place to start if one needs a sky model that easily integrates with modern MC rendering technology, we derived a suitably extended form of the Prague Sky Model [WVBR*21] as an initial reference model for these communities.

The new model is only hemispherical, i.e. limited to the ground-level observer altitude, as this is the main use case for photovoltaic simulation and building analysis. It retains all the other components of the Prague Sky Model, such as solar radiance patterns, a transmission term, and even polarisation patterns. The main difference is a considerably extended spectral range (280 nm to 2480 nm), and the corresponding adaptations that were necessary in the brute force pre-computation step. Outside the visible range, several additional factors need to be included in the simulation, which we discuss in Section 3.2. We also provide an implementation, both for evaluation of the new model in a renderer, but also an interactive visualisation tool for exploration of the dataset.

2. Related work

As outlined in the related work section of [WVBR*21] (to which we refer the reader for an exhaustive overview of related graphics literature), there are four main approaches to obtaining solar irradiance data in simulations:

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

For the spectral range we are investigating in this paper (UV + visible + near-IR), so far only option 0 is practical: and even that has only been achieved in software frameworks that are not directly conducive to the estimation of photovoltaic plant yield and building simulations, for reasons we discuss in the following paragraphs. Captured datasets (option 1) for such wide-band spectral ranges do not exist, and interactive approximations like [Hil20] (option 2) are problematic for our purpose due to the potentially unbounded error they introduce. Also, there is no data on how to make interactive

techniques perform reasonably well in the near-infrared. And finally, to our knowledge, no fitted model (option 3) that covers the spectral range needed for full photovoltaic assessments exists.

2.1. Brute Force Solvers

As for solvers that fall into category 0, the atmospheric research community has developed several highly interesting toolkits for hyper-spectral reference simulations of radiative transfer in the Earth’s atmosphere. LibRadtran [EBSK*16] offers excellent accuracy, a variety of solvers, and also reliably covers wavelengths far from the visible range. However, its accuracy comes at the cost of computation times that are too slow for practical use: this was the main reason to develop the Prague Sky Model, which in the visible range offers output similar to what libRadtran can compute, based on OPAC data [HKS98]. The fitted model is orders of magnitude faster to evaluate than the brute force computations and thereby becomes usable in a rendering context.

The SMARTS spectral irradiance model [Gue19] provides high-resolution spectral irradiance output for potentially sloped surfaces at ground level, and for a variety of clear-sky scenarios. Over 25+ years of usage, it has been extensively validated, and is widely used in a large number of simulation fields.

The remote sensing community has developed a number of wide-band hyper-spectral reference solvers, such as DART [GGEG*09]. Several other software packages of roughly the same type exist (e.g. 6S [VTD*97]): common to all of them is that hyper-spectral predictions for scene appearance *when viewed from orbit* are provided, to aid with the interpretation of real data acquired during satellite passes. In the case of DART, satisfactory performance is achieved via careful trade-offs between discretisation and other approximations. For the purposes it is intended for, the accuracy it reliably delivers is perfectly sufficient and has become a standard tool for such applications. However, for simulation of photovoltaic plant yield *viewed from the ground*, or for adaptation to a more general MC solver for building analysis, all these solvers suffer from the fact that they are custom-made for the specific purpose of “seeing things from orbit” and altering them to provide such functionality would mean a partial or complete rewrite.

Common to all these models is that they are very sophisticated, highly accurate and extensively verified but that they would be hard to integrate into modern MC rendering software. This was not their original purpose, so this fact in no way is to be held against them. But, as already stated in the introduction, light transport simulation technology is generally moving in the direction of MC rendering, due to its significantly higher predictive capabilities for complex scenes and materials.

2.2. Fitted Analytical Models

The Prague Sky Model [WVBR*21] served as the basis of our work, so we give a brief overview of it. As stated in the introduction to this section, we refer the reader to the related work section of that paper for an in-depth overview of other fitted models.

The Prague Sky Model follows the general approach of previous

fitted models by first running brute force simulations and then fitting a model to the obtained data. However, it advances the state of the art with regard to practically all components of this process. It uses standard OPAC clear sky data [HKS98] 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 spectral reference images. Then comes the fitting, which is essentially a smart compression scheme: through a projection to a suitable parameter space and the application of Canonical Polyadic Decomposition (CPD) the ≈ 250 GB of reference images are brought down to a dataset of ≈ 2.5 GB size. A matching transmission function computed from the same atmospheric data is provided, as well as first-order polarisation patterns.

Contrary to previous models such as the one by Perez [PSM93], the Prague Sky Model can be evaluated for observer altitudes greater than zero and observer altitudes up to 15 km are provided. The model is further parameterised by visibility (horizontal viewing distance at ground level), solar elevation, ground albedo, and wavelength channel. The spectral resolution is not very high, with 11 channels in the near-UV range and visible range (320 nm – 760 nm). However, as the target for this model are image synthesis computations in the visible range, mainly for consumption by human observers, such a comparatively sparse sampling of the spectral domain makes sense.

3. Near-infrared Extension

In this section, we present our main contribution: a hemispherical version of the Prague Sky Model [WVBR*21] with near-infrared wavelength channels. We followed the same steps as the original model: first generate the reference dataset using a path tracer, and then perform the fitting using CPD. For both steps, we used the original codebase of [WVBR*21]. However, we had to make several modifications to the codebase to achieve our goal, particularly to the dataset generation.

3.1. Spectral Range

The original Prague Sky Model provides 11 channels of 40 nm width, in the range from 320 nm to 760 nm. In order to cover the wavelength range up to 2500 nm we added 43 more such channels from 760 nm to 2480 nm. We also added one more such channel on the ultraviolet side of the spectrum, as there is some residual solar radiation as well: so our extended model covers wavelengths from 280 nm to 2480 nm using 55 regularly spaced bins. This is not a very high spectral resolution, at least not by the standards of atmospheric science: but as our path tracer splats MC samples into the spectral bins via a tent kernel, spectral aliasing is kept low, and the overall energy of the result spectra is maintained.

The exact terminology of what constitutes near-infrared and short-wavelength infrared (SWIR) varies by application area: the boundary between SWIR and medium-wavelength infrared is usually placed somewhere in the region between 3000 nm and 4000 nm. As we are focused on solar radiation, which at ground level practically goes to zero beyond 2500 nm, we limit the range of our model to near that value: explicitly covering the entirety

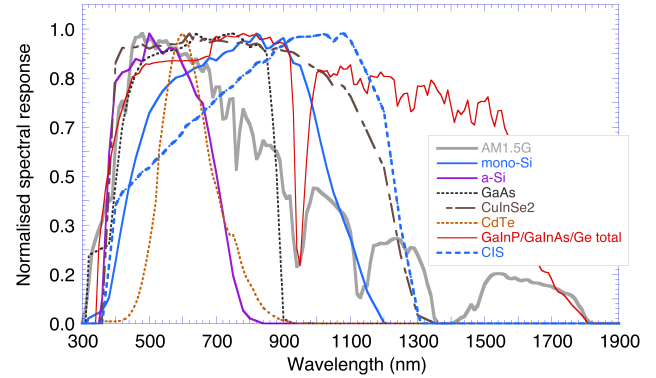


Figure 2: Spectral response curves for several solar cell types, with typical ground level solar radiation shown in grey. Note that several types absorb significant amounts of radiation way beyond the visible range: any predictions of photovoltaic plant yield based only on visible wavelengths would be considerably in error. Image courtesy of Chris Guyemard, used by permission.

of SWIR was not our goal, especially as there are no really official definitions for these region boundaries anyway. As Figure 2 shows, a model that explicitly only targets photovoltaic simulations could probably have stopped even earlier (around 2000 nm): but we wanted to cover the entire solar irradiance spectrum, and even though it does not have a lot of energy beyond 2000 nm, this might still matter for thermal irradiance analysis.

3.2. Dataset Generation

Since the original dataset generation method only works with a limited spectral range, it lacks some of the atmospheric data needed for the longer wavelengths we are interested in. Also, `atmo_sim`, the path tracer used by [WVBR*21], is based on ART [Wil18], which in its original form only supports UV plus the visible range, so we had to create a SWIR-capable version of that rendering toolkit. These changes were both to the internals of the toolkit, but also to the image formats used for saving the resulting images. Due to the additional output data we were generating being outside the visible range, we performed substantial validation of its output to make sure it is correct even in the previously untested wavelength ranges.

3.2.1. Input Data

The atmosphere of the Prague Sky Model follows data provided in the OPAC model [HKS98], and consists of several components: Rayleigh scattering on air molecules, absorption on ozone molecules, and Mie scattering and absorption on three types of aerosols – INSO, WASO and SOOT. Rendering of these components is governed by altitude-dependent particle concentration and wavelength-dependent scattering and absorption cross-section. Additionally, the Mie scattering also requires a wavelength-dependent phase function asymmetry parameter. While the concentrations are valid for any wavelength, the cross-sections and asymmetry parameters had to be extended to cover the entire 280 – 2480 nm range.

Rayleigh scattering cross-section for the additional wavelengths

was computed analytically using Bodhaine’s formula [BWDS99]. The ozone absorption cross-section was filled in with measurements of Gorshelev et al. [GSW*14]. Although they end at 1100 nm, they reach zero at that point, so the rest of the range could be neglected. Aerosol cross sections and asymmetry parameters were obtained from libRadtran [EBSK*16] and OPAC [HKS98].

Besides supplying missing data for components already present in the model, we also had to add entirely new components: absorption by CO₂, H₂O, and O₂ molecules. These were omitted from the original Prague Sky Model since their effect for wavelengths shorter than 760 nm is negligible. However, they are responsible for significant absorption bands in the infrared part of the spectrum, as shown in Figure 3. Vertical concentration profiles for these new absorbers came from the U.S. Standard Atmosphere [ACK*86], and their absorption cross-sections were extracted from the HITRAN database [GRH*22]. The new quantities are shown in Figure 4.

Finally, we also had to extend the used extraterrestrial solar radiance spectrum based on Wehrli’s measurements [Weh85].




Figure 3: Comparison of our atmospheric model to real sky dome measurements by Kider et al. [KKN*14]. Blue: Radiance spectrum from the Kider dataset for 27-05-13, 09:30, averaged over all provided sample points. Red: Radiance spectrum rendered by our path tracer for a similar atmospheric configuration, averaged over the same sample points. Black dashed: Radiance spectrum rendered by our path tracer but for an atmosphere without the newly added absorbers. Note how the additional absorbers yield much better match to the real data at wavelengths marked by the red arrows.

3.2.2. Path Tracer

After extending the input data, we had to also extend all data types used inside the path tracer for storing them. More importantly, we had to also modify the way the path tracer samples wavelengths and stores results so it can simulate the entire 280 – 2480 nm range and output an image with 55 channels. As expected, this significantly lowered the path tracer performance. To at least partially alleviate this slowdown, we increased the original number of 4 simultaneously traced Hero wavelength samples [WND*14] to 16. See

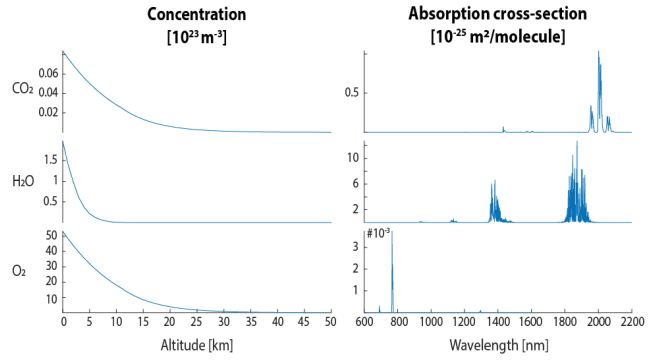


Figure 4: Vertical concentration profiles and absorption cross-sections of the newly added absorbers. The peaks in cross-sections are responsible for the absorption bands visible in Figure 3.

Section 5.2 for a discussion of the resulting performance: however, please note that the worse performance of the SWIR-capable path tracer only affects the pre-computation step. Also, as discussed in Section 5.2, a sizeable part of this slowdown is intrinsic to the problem at hand (far higher atmospheric absorption in certain wavelength bands outside the visible range), so that the much lower rendering speeds for a SWIR solution are actually a good argument to use a pre-computed model like ours in the first place.

Finally, we double-checked the used scattering phase functions and found that there are modifications needed there as well. The original path tracer approximates Mie phase functions using Henyey-Greenstein functions [HG41]. While this approximation is sufficient in the original wavelength range, for SOOT aerosols it starts to deviate significantly from the true Mie phase function for longer wavelengths. The Cornette-Shanks phase function [CS92] is a much better match in this range, as shown in Figure 5. However, since the contribution of the SOOT aerosol is small, we found the overall impact of this change to be negligible.

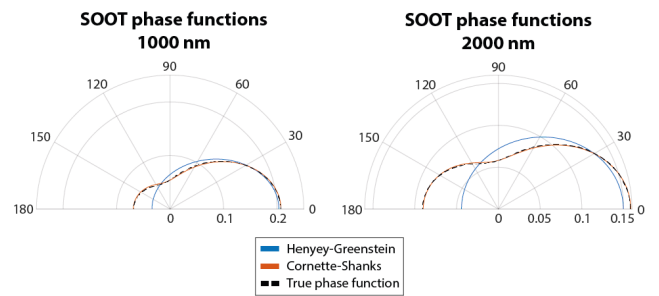


Figure 5: Comparison of Henyey-Greenstein (blue) and Cornette-Shanks (red) phase functions with the true Mie phase function (black dashed) provided for the SOOT aerosol by OPAC [HKS98]. While Cornette-Shanks matches the true phase function well, Henyey-Greenstein deviates significantly for longer wavelengths.

3.2.3. Validation

In order to make sure that our modified dataset generation process yields correct output, especially for the infrared wavelengths, we

validated it against real sky dome measurements by Kider et al. [KKN*14]. They provide a systematically collected dataset which includes spectral sky dome radiance measurements taken in the course of several days for 81 sample positions in the sky.

We selected one particular date and time (27-05-13, 09:30) and computed our reference spectrum by averaging the corresponding 81 measured spectra. Then we tried to match this reference with our modified renderer. The only drawback of work of Kider et al. is that it does not contain the exact atmospheric parameters at the time of capture so a perfect match is not easily possible. We got the best match for visibility 59.4 km (which corresponds to OPAC continental average aerosol composition), solar elevation 41.83° and ground albedo 0. Our result is compared with the reference spectrum in Figure 3. Given the unknown atmosphere composition, we consider the match very good. Especially the absorption bands match nicely.

3.3. Fitting

After generating the dataset, we moved to the next step in model creation – fitting. In the Prague Sky Model the fitting corresponds to compressing the dataset using CPD, i.e. each reference image is modelled as a sum of several outer products of vector pairs, and the fitting looks for vectors giving the best reconstruction. Since the original code fitted every wavelength channel of a reference image independently, modifying it to work with our extended range was straightforward. The only issue we encountered were artefacts caused by higher levels of noise in the reference images. As we discuss in Section 5.1, some wavelengths are more difficult to render because less light is transported in these spectral bands. As a result, there is more noise in the corresponding images, which makes the fitting less stable. To avoid undesirable artefacts, we had to increase the strength of the filtering step that is applied before the fitting in these cases. Figure 6 shows an example in which this modification successfully removed artefacts from the resulting fit.




Figure 6: An example of a successfully removed noise from a fitted image thanks to our slightly increased filtering in the fitting preprocessing step. Left: a reference image rendered for visibility 59.4 km, solar elevation 0°, ground albedo 0.5 and wavelength channel 2400 – 2440 nm. Middle: a fit obtained with the original code of Prague Sky Model. Right: a fit after increasing the filtering.

3.4. Wavelength Compression

Using 5× more wavelength channels than the Prague Sky Model does not come for free. Besides slowing down the generation of the reference dataset and its fitting, it also means a 5× larger final dataset that constitutes the model. This dataset has to be stored on

disk and then loaded into computer memory before use. Therefore, we also investigated the possibility of compressing the dataset by omitting some of the 55 channels we computed.

We tested a simple greedy algorithm for selecting unnecessary channels. We kept the first 12 channels (280 – 760 nm) fixed and then tried leaving out immediately following channels as long as the maximum relative error caused by replacing them by interpolated values does not exceed a fixed threshold. Once that happens, the last tested wavelength is kept and the omitting starts from the next one. Figure 7 shows dataset size reduction we can achieve with this algorithm for different error thresholds. So if a particular application does not require maximum accuracy and can allow e.g. 15% error, it can save almost 22% of the dataset size.

To ensure maximum quality and flexibility at the same time, we supply the complete dataset with all channels, together with a list of channels that can be omitted for various error thresholds.

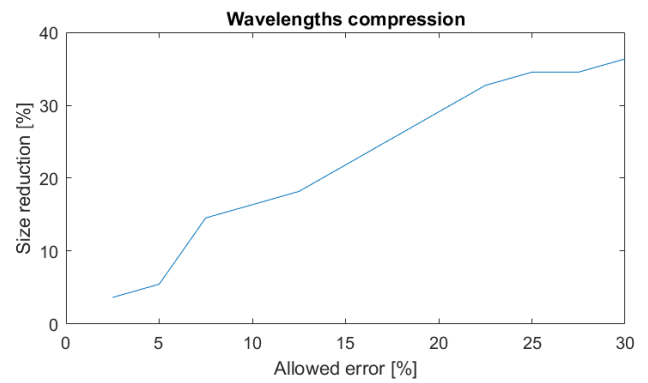


Figure 7: Size reduction of the final dataset that is achievable by replacing some of the 55 channels with interpolated values while keeping the introduced error under a selected threshold.

4. Implementation

Our secondary contribution is a standalone implementation of the model. While Wilkie et al. provided an implementation with the Prague Sky Model, it was a sub-optimal research code written in C and Objective-C that was difficult to use. Our goal was to create a clean and optimised C++ 17 implementation that would be sufficiently documented and easy to use. It consists of 3 parts: model library, example renderer, and front end.

The library is formed by just two files and everything one needs to use the sky model is in this code. It works with our new dataset as well as with every dataset previously released for the Prague Sky Model. In comparison with the original implementation, our code significantly lowers memory consumption. Instead of loading the entire dataset into memory and completely unpacking it there, we perform part of the unpacking on demand and provide an option to load only a part of the dataset needed for rendering a selected configuration. This way we can reduce memory use up to 24 times.

To illustrate how to use the library, we implemented a simple example renderer. It shows how to query the library to render a fisheye

image of the sky. On AMD Ryzen 9 3900X 12-core processor one 1000x1000 pixels large image with all 55 channels takes around 600 ms to render which allows use in interactive applications.

Finally, for immediate model testing and dataset exploration, we accompany the example renderer with front end with both command-line and graphical user interfaces (GUI). Using the GUI (shown in Figure 8) users can easily load datasets, interactively change rendered configurations, and save rendered images.

Our implementation was tested on Windows and Linux and will be released under the Apache-2.0 licence.

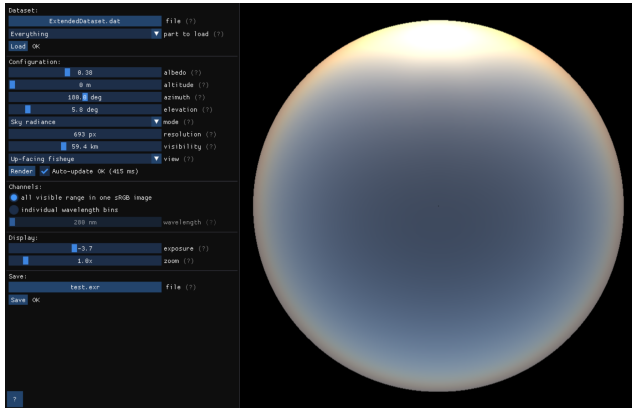


Figure 8: Screenshot of GUI provided by our implementation.

5. Results and Limitations

Our model provides sky radiance, polarisation, and transmittance for zero observer altitude and 55 regularly spaced wavelength channels from 280 nm to 2480 nm. The remaining model parameters, i.e. visibility, solar elevation, and ground albedo, are the same as in the Prague Sky Model. A sample of this output is shown in Figure 9. For example, the first channel (280 – 320 nm) shows that although the transmittance is low due to a strong ozone absorption, the sky radiance is relatively high compared to the other channels because of the high extraterrestrial solar radiance there. On the other hand, low solar radiance makes the sky radiance in the last channel (2440 – 2480 nm) very low even though the transmittance is high. And when the decreasing solar radiance meets with the low transmittance due to the water absorption in channel 1360 – 1400 nm, the resulting sky radiance is almost zero. Another different set of channels and model parameters is shown in Figure 1.

5.1. Fitting error

An important aspect of every fitted model is the amount of error introduced by the fitting. Here we discuss errors with respect to the wavelength channels as the behaviour with respect to the visibility, solar elevation and ground albedo stays the same as in the Prague Sky Model. Figure 10 shows a box plot of the normalised mean absolute errors between the fits and reference images grouped by the wavelength channels. For most channels, the median of the errors stays under 10%, the total average error is 6%. However, there are

three cases where the error is significantly larger: in the near-UV channel 280 – 320 nm, between 1360 and 1400 nm, and between 1840 and 1960 nm. All three cases correspond to strong absorption bands caused by ozone and water, which make atmosphere simulation by the path tracer more difficult. More traced paths end up being absorbed, and by the time other channels have already converged, the noise level in these three bands remains high. As a result, the error between the smooth fit and noisy reference image increases. However, as the grey overlay in the figure shows, these regions carry very low energy, making them essentially irrelevant. Similarly, all channels above 2000 nm exhibit larger errors than the channels in the visible range but the unnormalised mean absolute error is practically zero there. Figure 11 gives an example of a configuration with the average error.

5.2. Limitations

In terms of practical usability of our model, the biggest limitation is the size of the final dataset. The hemispherical version presented in this paper, with its extended set of wavelength channels, requires 550 MB of disk space and twice as much computer memory when loaded for use. Using the same compression technology, a full altitude-resolved model with the same extended channels would consume 12.1 GB of disk space and 24.2 GB of computer memory: however, it has to be noted that usage cases for altitude-resolved SWIR sky models seem to be scarce. We also propose a method for omitting unnecessary channels, which can reduce the dataset size if some additional error can be tolerated in a given application.

As far as further development of the model in its current form is concerned, we see the main limitation in the resource-intensive reference dataset generation. This means rendering a large number of images to a low level of noise, which requires a lot of time and computational resources. Rendering of our hemispherical extended dataset consumed more than 800 thousand core-hours and was done in the course of one week on a scientific supercomputing cluster.

With SWIR added, the rendering step of the pre-computation runs significantly slower than in the Prague Sky Model because of the additional wavelengths. First, there is strong absorption in some of the new wavelength bands, which makes rendering in these bands considerably more noisy. Second, the necessary modifications of the path tracer slowed it down almost 20 times. A factor of 5 comes from a more complex atmosphere description and the overhead of rendering 55 channels in one image. And a factor of 4 comes from twice larger variance for the same number of samples. We raised the original number of 4 simultaneously traced Hero wavelength samples to 16, but this gave us only a speedup by a factor of two. However, the slowdown affects the pre-computation step only. Rendering a single wavelength channel using our model takes the same amount of time as with the Prague Sky Model.

6. Conclusion and Future Work

We presented a pre-computed model of sky radiance, transmission and polarisation for a much wider spectral range than such models offered before. We cover practically the entire range of solar irradiance at ground level, which makes our new model useful for predictive simulations of photovoltaic plant yield and thermal building

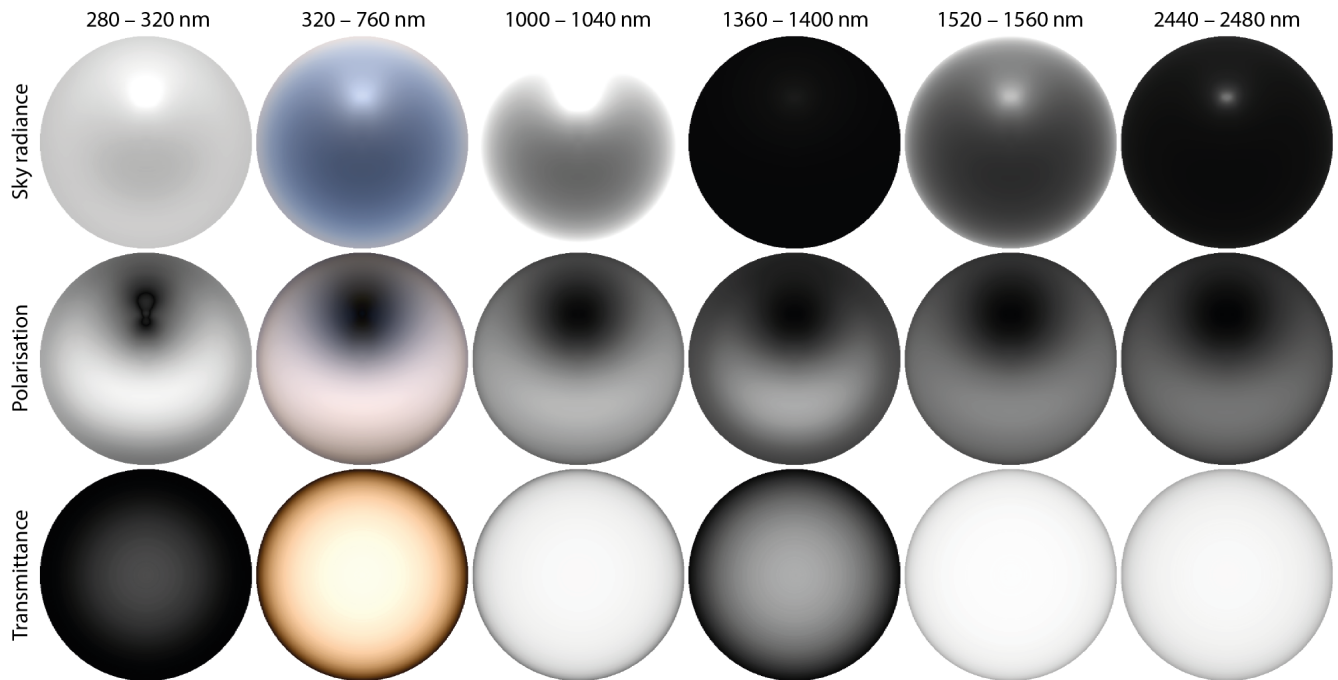


Figure 9: An example of outputs provided by our model: sky radiance, polarisation and transmittance for six wavelength bands. All images show up-facing fisheye view of the sky with the same configuration of visibility 59.4 km, solar elevation 45° and ground albedo 0.5. The second column was obtained by combining corresponding channels into an sRGB image. Each row was tone-mapped independently, but all images in one row share the same exposure, except for the second column.

assessments. We validated the additions we had to make to the path tracing step to properly cover the infrared range, so that the new model can serve as reference for spectral path tracing in the entire wavelength bracket of solar irradiance.

We also created a better retrieval implementation of the model than the Prague Sky Model originally provided, along with an interactive visualisation and exploration tool for the dataset.

Future work will go in the direction of optimising the used path tracer, better wavelength compression, and rendering more atmospheric configurations.

7. Acknowledgements

We acknowledge funding by Chaos Czech, the Czech Science Foundation grant 22-22875S, the EU H2020-MSCA-ITN-2020 grant number 956585 (PRIME), and the Charles University grant number SVV-260588. This work was supported by the Intel Graphics and Visualization Institute at Charles University. The authors also acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing HPC resources.

References

- [ACK*86] ANDERSON G., CLOUGH S., KNEIZYS F., CHETWYND J., SHETTLE E.: AFGL atmospheric constituent profiles (0.120km). 46. 4
- [BWDS99] BODHAINE B. A., WOOD N. B., DUTTON E. G., SLUSSER J. R.: On Rayleigh optical depth calculations. *Journal of Atmospheric and Oceanic Technology* 16, 11 (1999), 1854–1861. 4
- [CS92] CORNETTE W. M., SHANKS J. G.: Physically reasonable analytic expression for the single-scattering phase function. *Appl. Opt.* 31, 16 (Jun 1992), 3152–3160. 4
- [EBSK*16] EMDE C., BURAS-SCHNELL R., KYLLING A., MAYER B., GASTEIGER J., HAMANN U., KYLLING J., RICHTER B., PAUSE C., DOWLING T., BUGLIARO L.: The libRadtran software package for radiative transfer calculations (version 2.0.1). *Geoscientific Model Development* 9, 5 (2016), 1647–1672. 2, 4
- [GGEG*09] GRAU E., GASTELLU-ETCHEGORRY J.-P., GASCON F., RUBIO J., BRUT A.: Earth-atmosphere radiative transfer in DART model. In *WHISPERS* (2009), IEEE, pp. 1–4. 2
- [GRH*22] GORDON I., ROTHMAN L., HARGREAVES R., HASHEMI R., KARLOVETS E., SKINNER F., CONWAY E., HILL C., KOCHANOV R., TAN Y., WCISŁO P., FINENKO A., NELSON K., BERNATH P., BIRK M., BOUDON V., CAMPARGUE A., CHANCE K., COUSTENIS A., DROUIN B., FLAUD J., GAMACHE R., HODGES J., JACQUEMART D., MLAWER E., NIKITIN A., PEREVALOV V., ROTGER M., TENNYSON J., TOON G., TRAN H., TYUTEREV V., ADKINS E., BAKER A., BARBE A., CANÈ E., CSÁSZÁR A., DUDARYONOK A., EGOROV O., FLEISHER A., FLEURBAEY H., FOLTYNOWICZ A., FURTENBACHER T., HARRISON J., HARTMANN J., HORNEMAN V., HUANG X., KARMAN T., KARNS J., KASSI S., KLEINER I., KOFMAN V., KWABIA-TCHANA F., LAVRENTIEVA N., LEE T., LONG D., LUKASHEVSKAYA A., LYULIN O., MAKHNEV V., MATT W., MASSIE S., MELOSSO M., MIKHAILENKO S., MONDELAIN D., MÜLLER H., NAUMENKO O., PERRIN A., POLYANSKY O., RADDAOUI E., RASTON P., REED Z., REY M., RICHARD C., TÓBIÁS R., SADIEK I., SCHWENKE D., STARIKOVA E., SUNG K., TAMASSIA F., TASHKUN S., VANDER AUWERA J., VASILENKO I., VIGASIN A., VILLANUEVA G., VISPOEL B., WAGNER G., YACHMENEV A., YURCHENKO S.: The HITRAN2020 molecular spectroscopic database. *Journal of Quantitative Spectroscopy and Radiative Transfer* 277 (2022). 4




Figure 10: A box plot of the mean absolute error between fits and their corresponding reference images, normalised by the average value of each channel, shown for all channels in our model. The red line is the median, the blue box goes from the first to the third quartile, and the whiskers are the minimum and maximum values. The increased error around 280, 1360 and 1840 nm is due to noise in the reference images caused by strong absorption in these bands. The grey shape in the background shows the normalised average pixel value in each channel.




Figure 11: An example of configuration with the average fitting error. It corresponds to visibility 59.4 km, solar elevation 18.64°, ground albedo 0 and wavelength channel 1320–1360 nm, the normalised mean absolute error is 0.059. The error is barely noticeable in the difference image and completely invisible in the fit.

- [GSW*14] GORSHELEV V., SERDYUCHENKO A., WEBER M., CHEHADE W., BURROWS J. P.: 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. 4
- [Gue19] GUEYMARD C. A.: The SMARTS spectral irradiance model after 25 years: New developments and validation of reference spectra. *Solar Energy* 187 (2019), 233–253. 2
- [HG41] HENYEV L. G., GREENSTEIN J. L.: Diffuse radiation in the Galaxy. *Astrophysical Journal* 93 (Jan 1941), 70–83. 4
- [Hil20] HILLAIRE S.: A scalable and production ready sky and atmosphere rendering technique. *Computer Graphics Forum* 39, 4 (2020),

13–22. 2

- [HKS98] HESS M., KOEPKE P., SCHULT I.: Optical properties of aerosols and clouds: The software package OPAC. *Bulletin of the American Meteorological Society* 79, 5 (1998), 831–844. 2, 3, 4
- [KKN*14] KIDER JR. J. T., KNOWLTON D., NEWLIN J., LI Y. K., GREENBERG D. P.: A framework for the experimental comparison of solar and skydome illumination. *ACM Trans. Graph.* 33, 6 (Nov. 2014), 180:1–180:12. 4, 5
- [PSM93] PEREZ R., SEALS R., MICHALSKY J.: All-weather model for sky luminance distribution—preliminary configuration and validation. *Solar Energy* 50, 3 (1993), 235–245. 1, 3
- [VTD*97] VERMOTE E., TANRE D., DEUZE J., HERMAN M., MORCETTE J.-J.: Second simulation of the satellite signal in the solar spectrum, 6S: an overview. *IEEE Transactions on Geoscience and Remote Sensing* 35, 3 (1997), 675–686. 2
- [Weh85] WEHRLI C.: Extraterrestrial solar spectrum. In *Publication no. 615*. Physikalisch-Meteorologisches Observatorium + World Radiation Center, Davos Dorf, Switzerland, July 1985. 4
- [Wil18] WILKIE A.: The Advanced Rendering Toolkit, 2018. <http://egg.mff.cuni.cz/ART>. 3
- [WND*14] WILKIE A., NAWAZ S., DROSKE M., WEIDLICH A., HANIKA J.: Hero wavelength spectral sampling. In *Computer Graphics Forum* (2014), vol. 33, Wiley Online Library, pp. 123–131. 4
- [WVBR*21] WILKIE A., VÉVODA P., BASHFORD-ROGERS T., HOŠEK L., ISER T., KOLÁŘOVÁ M., RITTIG T., KRIVÁNEK J.: A fitted radiance and attenuation model for realistic atmospheres. *ACM Trans. Graph.* 40, 4 (July 2021). 1, 2, 3