Eurographics Symposium on Rendering (DL-only Track) (2021)
A. Bousseau and M. McGuire (Editors)

Sampling Clear Sky Models using Truncated Gaussian Mixtures

N. Vitsas[1] and K. Vardis[1] and G. Papaioannou[1]

[1]Department of Informatics, Athens University of Economics and Business, Greece

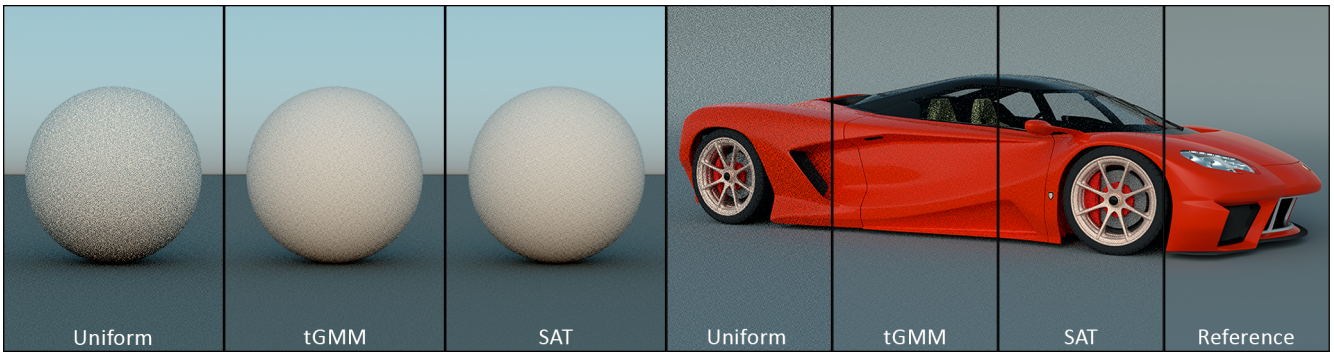

Figure 1: *Two example scenes demonstrating that our radiance-based model (tGMM) is able to distribute samples towards the important regions of the sky, offering variance reduction a) significantly higher than naive sampling (uniform) and b) that matches more computationally demanding approaches (SAT).*

Abstract

Parametric clear sky models are often represented by simple analytic expressions that can efficiently generate plausible, natural radiance maps of the sky, taking into account expensive and hard to simulate atmospheric phenomena. In this work, we show how such models can be complemented by an equally simple, elegant and generic analytic continuous probability density function (PDF) that provides a very good approximation to the radiance-based distribution of the sky. We describe a fitting process that is used to properly parameterise a truncated Gaussian mixture model, which allows for exact, constant-time and minimal-memory sampling and evaluation of this PDF, without rejection sampling, an important property for practical applications in offline and real-time rendering. We present experiments in a standard importance sampling framework that showcase variance reduction approaching that of a more expensive inversion sampling method using Summed Area Tables.

CCS Concepts

• ***Computing methodologies*** → ***Ray tracing; Image-based rendering; Mixture modeling;***

1. Introduction

Analytic sky models have been an active area of research for many years [NDKY96, PSS99, HW12, Hil20]. This is mainly due to their applicability in a wide range of rendering tasks, from real time rendering to architectural simulations. In many environments, natural light from the sky dome can be a dominant factor to a scene's illumination, especially when the sun disk is not directly visible. These parametric models have the benefit of adjustable sun position and clearness conditions through a small and configurable set of parameters, making them intuitive to adjust and artist-friendly. Even though a lot of work has focused on the generation of analytic formulas for the radiance distribution across the sky, the efficient sampling of such distributions remains an open problem. Typically, analytic sky models are treated as regular environment maps that are sampled either uniformly or using a discrete distribution on some parametric space.

In this work, we are particularly interested in achieving efficient sampling of the sky illumination directly from the parametric sky model, potentially also spread over time and conditions, as is the case in architectural simulations and daylighting systems (Figure 2). As stochastic methods based on ray tracing become more practical thanks to recent hardware advances, an analytic formula for sampling analytic sky models becomes very useful. Most modern ray-tracing-based rendering approaches for outgoing radiance

DOI: 10.2312/sr.20211288

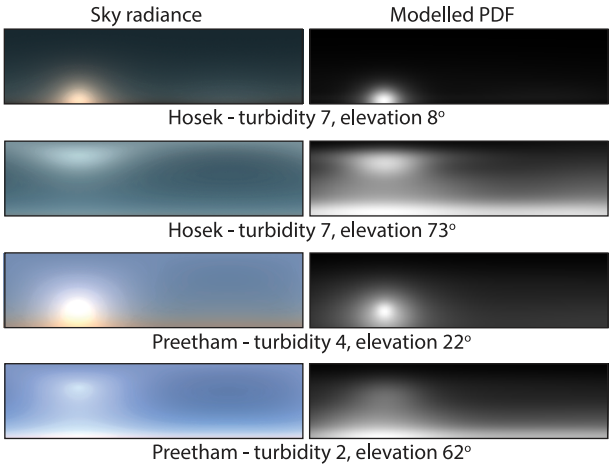

Figure 2: *Sampling of different sky model instances using our tGMM-based PDF model. Left: sky radiance in hemispherical coordinates. Right: resulting PDF, accounting for the polar distortion.*

estimation rely on Monte Carlo integration with importance sampling. A distribution function for this sampling process needs to be efficient to compute and straightforward to incorporate in any such framework [Vea98, TCE05]. One potential direction would be to work directly on the analytic formula provided by the chosen sky model. However, this can be computationally expensive or mathematically infeasible. Furthermore, this would restrict the approach to a particular sky model. Instead, we decided to fit a new distribution to the underlying analytic sky model's radiance, a *truncated Gaussian mixture model* (tGMM) in particular, adjusting its parameters to match the specific sky model in question.

Mixture models are widely used in data mining, pattern recognition, machine learning, and statistical analysis [Bis06]. Apart from providing a framework for clustering data, they can also be used for building complex probability distributions. A very common and well-behaved mixture model is the Gaussian mixture model (GMM), which is a linear combination of Gaussian components, aiming to provide a richer class of density models than a single Gaussian. A tGMM is simply a linear combination of truncated Gaussian components. Truncation refers to bounding the distribution to a region $[\alpha,\beta], \alpha,\beta \in \mathbb{R}$, outside of which the PDF is zero. The small storage requirements of a single tGMM allows us to rapidly generate and efficiently store a large number of tGMMs that compactly encode the radiance-based distribution of an entire sky model.

Our work is based on two key observations: a) the radiance distribution of the sky exhibits low frequency changes and uniform regions and, thus, can be well-fitted with normal distribution kernels, and b) bounding, i.e. truncating, the distribution to the hemispherical domain dispenses with (expensive) rejection sampling approaches, which would be inevitably required at least for the latitudinal sample coordinate, due to the hard clipping at the horizon. Our experiments show that a tGMM of only a few components can accurately approximate the distribution of the sky model and is very easy to store, sample and evaluate.

To summarise, the main contribution of this paper is a probability density function based on truncated Gaussian mixture models fitted on the sky radiance distribution that:

- uses minimal and adjustable look-up-table (LUT) storage to encode the approximate PDF,
- has constant-time evaluation, offering great performance characteristics for both CPU and GPU implementations,
- can easily be integrated in any existing importance sampling estimator,
- dispenses with any form of rejection sampling, which is unpredictable in terms of performance.

A brief overview of the proposed model generation and sky model importance sampling is shown in Figure 3.

2. Related Work

This section covers previous work on analytic sky models, environment map sampling techniques and the application of learned distributions, like GMMs, to several aspects of rendering.

Analytic sky models. The importance of using realistic and controllable sky illumination in rendering has led to several parametric methods being proposed over the years. Some of them, have been formally adopted by CIE [C*94]. Many well-known and widely used models in computer graphics are based on variations of the Perez formula [PSM93]:

$$\mathcal{F}(\theta,\gamma) = (1 + Ae^{B/cos\theta})(1 + Ce^{D\gamma} + Ecos^2\gamma). \quad (1)$$

The two input variables are θ and γ, which are the angle between the zenith and the viewing direction and the angle between the sun and the viewing direction, respectively. The terms A, B, C, D, E have specific physical effects on the sky's distribution like darkening and brightening of the horizon or relative intensity of the circumsolar region.

Based on the Perez formula, Preetham et al. [PSS99] proposed a fitting of the parameters using reference simulated data. Hosek et al. [HW12] expanded the Preetham model using data generated by an offline volumetric path tracer, taking into account complex atmospheric phenomena. Hosek's model is more stable and offers a wider range of sky appearances. The main configurable parameter in both models, apart from the sun position (altitude, longitude) is turbidity t, which models the clearness of the sky. In the Preetham model the valid range for t is $[2,6]$, while Hosek performs well in a wider range, $[2,10]$. One can rapidly generate sky maps by manipulating those few parameters. The Hosek model also adds an *albedo* parameter ρ, which mainly affects the chromatic appearance of the sky rather than its power distribution and therefore, we exclude this parameter from our tabular PDF data.

An even more expressive and compact sky and atmospheric model was recently proposed by Hillaire [Hil20]. This model is very efficient and can handle atmospheric visualisation from any point of view, ground or space with focus on real-time graphics and artistic visualisation. A machine-learning-driven process for learning and predicting sky models was proposed by Satilmis et

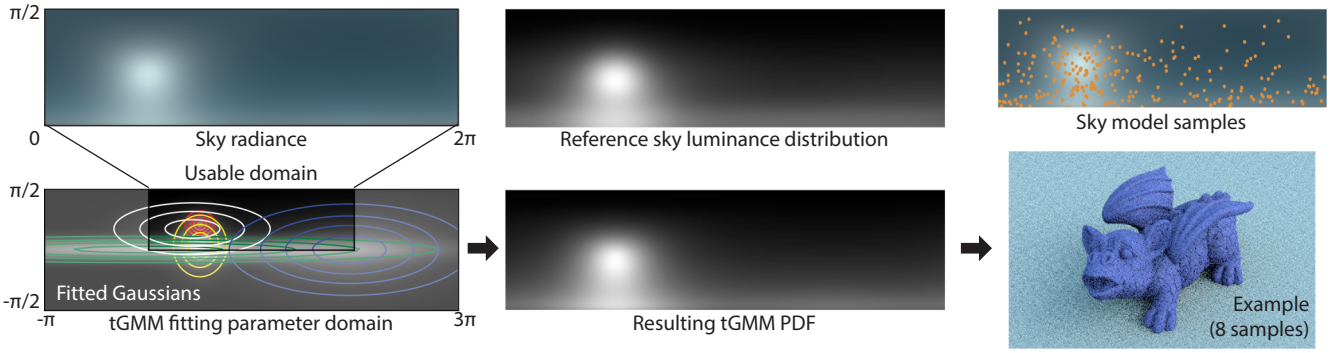

Figure 3: *The fitting and sampling process of our analytic PDF, applied to the Hosek sky model.*

al. [SBCD17]. Using a single layer artificial neural network, the authors show that the network can adapt to several analytic models and even captured images. However, it is not discussed how this can be integrated in an importance sampling framework. The interested reader can refer to the work of Bruneton [Bru17] for an extensive evaluation of several clear sky models.

Compact sky radiance representations. Efficient encoding and storage of those sky maps can be achieved using long-standing representations in computer graphics. When approximations are acceptable, (Spherical) Radial Basis Functions have been extremely useful. They offer a very good compromise between storage, reconstruction performance and accuracy. Spherical gaussians [TS06], spherical harmonics [SKS02] and Haar wavelets [NRH03] have all been employed for fast, approximate precomputed light transport. Spherical harmonics in particular have been proposed for the Preetham model [HMW08].

Environment map sampling. Accurate sampling of environment maps is very important for image-based lighting. Sampling the sky dome naturally falls under the general category of environment sampling. A robust sampling approach is to use a cumulative density function, generated from a radiance-based PDF. Pharr et al. [PJH16] define a piecewise-constant 2D probability distribution function in image coordinates that corresponds to the distribution of radiance represented by the environment map. This function can then be sampled using the marginal and conditional CDF stored in two Summed Area Tables (SAT) [Cro84, CETC06, CAM08]. Previous research in environment sampling has also led to hierarchical subdivisions like Penrose tiling [ODJ04], stratified regions [ARBJ03] and "cascading sets" of conditional and marginal 1D CDFs [LRR05]. Other applicable approximate methods include spherical harmonics [JCJ09, Ber11] and wavelet importance sampling [CJAMJ05]. To account for occlusion that is most notably present in indoor scenes, Bitterli et al. [BNJ15] proposed visibility-based importance sampling of environment map regions through rectangular portals using a novel rectilinear mapping.

Several of the referenced methods for environment map sampling can accommodate specialised implementations for the specific problem of skylight sampling. However, our main concern is the construction of an analytic model that is simple, robust, easy to implement and can be sampled exactly.

Learning-based methods. The idea of learning a distribution from data has been explored heavily in the context of rendering, especially for approximating probability functions based on the product of direct lighting with the visibility or BRDF terms. In the context of path guiding, Vorba et al. [VKv*14] proposed guiding the sampling process by progressively learning sampling distributions based on GMMs in order to handle scenes with complex geometry and lighting conditions. Herholz et al. [HEV*16] extended this idea by fitting separately obtained GMMs for the BRDF and illumination terms, which are combined in a product mixture to importance sample indirect illumination. In a similar fashion, Tsai and Shih [TCJW08] used SRBF kernels to fit the BRDF and environment map which are then combined for efficient product sampling of the two. Reibold et al. [RHJD18] use complete transport paths of high variance as samples to construct and cache a continuous PDF using a Gaussian mixture model. These can later be sampled for important paths. Recently, Xia et al. [XWHM20] employed low-component Gaussian mixtures to approximate the BSDFs of complex layered materials. They take advantage of the analytic solutions to Gaussian products to importance sample multiple layers of a material.

3. Background

In this section, we review importance sampling, environment map sampling using Summed Area Tables and truncated Gaussian Mixture Models.

Importance sampling (IS) is a very effective variance reduction technique that has been applied to all terms of the rendering equation. It requires generating samples from probability density functions that exhibit similar characteristics to the original function. These samples can be generated by CDF inversion, rejection sampling, Metropolis sampling, or other means. When applicable, the inversion method is the most convenient and is the one used by most common BRDF and light source sampling techniques. However, many lighting distributions in a 3D scene are hard to sample using the inversion method. Ideally, in our sampling strategy, we want $p(x) \propto f(x)$, but a useful property of IS is that difficult or impossible to sample densities p can be replaced by an approximate distribution $\hat{p}$; better approximations lead to more efficient importance sampling. The result will remain unbiased as long as $\hat{p}$

integrates to 1 and is greater than zero wherever there is non-zero importance. Given the Monte Carlo estimator

$$I \approx \langle F_N \rangle = \frac{1}{N} \sum_i^N \frac{f(x)}{p(x)},$$

I is expected to converge faster, if samples are taken from a distribution that is similar to the integrand.

Sampling using Summed Area Tables. The standard and arguably more general and robust approach for importance-sampling an environment map is to consider a discrete 2D distribution with an unnormalised, piecewise constant PDF constructed from the intensity values I of each pixel [PJH16]. From this, one can easily calculate the marginal density function $p(v)$ and the conditional distribution $p(u|v)$ for each row of the image. During sampling, the marginal 1D distribution is used to select a row of the image to sample. Rows with bright pixels are more likely to be sampled. Consequently, the row is sampled using its 1D conditional distribution. The marginal and conditional CDFs are built from the discrete PDF values and stored in Summed Area Tables [Cro84]. However, the process of inversion sampling of these CDFs, although simple, requires an $O(\log n)$ operation in the form of a binary search over the SAT. Apart from the performance considerations, the calculation and storage of large tabular data for both the original image, as well as the CDFs and PDFs, is required. In execution environments like those of a GPU processor, such an operation can be unjustifiably costly.

Truncated Gaussian Mixture Model. To better understand the reasons we chose this representation as our model, we start with a brief discussion on Gaussian distribution components. In order to apply the inversion method to a distribution $p(x)$ we must first build its CDF $P(x)$ and its inverse $P^{-1}(x)$. The n-dimensional normal distribution takes the general form:

$$\phi(\boldsymbol{x};\boldsymbol{\mu},\Sigma) = \frac{\exp\left(-\frac{1}{2}(\boldsymbol{x}-\boldsymbol{\mu})^T\Sigma^{-1}(\boldsymbol{x}-\boldsymbol{\mu})\right)}{2\pi^{n/2}|\Sigma|^{1/2}}.$$

Sampling from this distribution with a non-diagonal covariance matrix Σ and arbitrary truncation regions would involve sample rejection and in our case, it would be unnecessary; due to the parameterisation of the sky domain and the inherent, symmetrical appearance of clear skies, we need only use a diagonal covariance matrix $\Sigma = diag(\sigma_1^2, \sigma_2^2)$, where the equation in 2D can be simplified to:

$$\begin{aligned}\phi(x,y;\mu_X,\mu_Y,\sigma_X^2,\sigma_Y^2) &= \frac{\exp\left(-\frac{1}{2}\left(\frac{x-\mu_X}{\sigma_X}+\frac{y-\mu_Y}{\sigma_Y}\right)\right)}{2\pi\sigma_X\sigma_Y} \\ &= \frac{\exp\left(-\frac{x-\mu_X}{2\sigma_X}\right)}{\sqrt{2\pi\sigma_X^2}}\frac{\exp\left(-\frac{y-\mu_Y}{2\sigma_Y}\right)}{\sqrt{2\pi\sigma_Y^2}} \\ &= \phi(x;\mu_X,\sigma_X^2)\phi(y;\mu_Y,\sigma_Y^2).\end{aligned}$$

This is the product of two 1D Gaussian distibutions, which can be sampled independently.

A truncated Gaussian in a region $[\alpha,\beta]$ can be exclusively formalised in terms of a regular Gaussian distribution:

$$\psi(x;\mu,\sigma^2,\alpha,\beta) = \frac{\phi(x;\mu,\sigma^2)}{\Phi(\alpha)-\Phi(\beta)}, \tag{2}$$

$$\Psi^{-1}(\xi) = \Phi^{-1}(\Phi(\alpha)+\xi(\Phi(\beta)-\Phi(\alpha)))\sigma+\mu,$$

where $\Phi(x;\mu,\sigma^2)$ and $\Phi^{-1}(x;\mu,\sigma^2)$ are the CDF and the quantile function of the regular Gaussian distribution, respectively. The above formulas demonstrate a useful property of truncated Gaussians; they can be defined and evaluated in terms of their "parent" Gaussian distribution. This effectively means that a truncated Gaussian can be sampled exactly, within the truncation bounds, using the analytic formulas of the unbounded Gaussian. Inverting a Gaussian function directly is not possible. It requires the use of the error function which, however, has very fast and robust analytic approximations that can be used. Moreover, the Gaussian distribution has been extensively studied and there are many methods to efficiently draw samples from it and its truncated variants.

A truncated Gaussian Mixture Model, similar to a Gaussian Mixture Model, can be expressed in terms of a latent selection variable $z \in [1..N]$:

$$p(x) = \sum_z p(z)p(x|z) = \sum_k^N \pi_\kappa \psi(x|\mu_k,\sigma_k), \tag{3}$$

where $0 \leq \pi_k \leq 1$ and $\sum_k^N \pi_k = 1$.

Sampling from such a mixture model is fairly straightforward, which makes it very useful in the context of importance sampling for rendering. The first step is to sample from the categorical distribution $p(z)$ defined by the weights of each individual component and then, draw a sample from the distribution of this one single component $\psi(x|\mu_k,\sigma_k)$.

4. Method Description

We describe here the complete process for the modelling, fitting and sampling of our tGMM PDF generated from the radiance distribution map of an analytic sky model. The method must generate samples over the polar coordinates of the sky dome hemisphere $\phi \in [0,2\pi), \theta \in [0,\frac{\pi}{2}]$, where θ is the sample elevation and ϕ its azimuth, that follow the distribution of the sky radiance (typically the luminance component). Unfortunately, drawing samples from the Perez model (Equation 1) with inversion sampling is very hard, if possible at all. However, the several exponential components in the equation and the resulting power distribution both hint at a reasonable fitting by a mixture model of Gaussian components, which we exploit.

4.1. PDF Modelling

We parameterise the PDF according to solar elevation and turbidity and obtain a tGMM for each parameter pair, by fitting the model to a normalised luminance map, densely generated from the sky model. The azimuth of the sun is not taken into account as a parameter for the model as it is typically only applied as a constant offset to the generated angle ϕ of each sample.

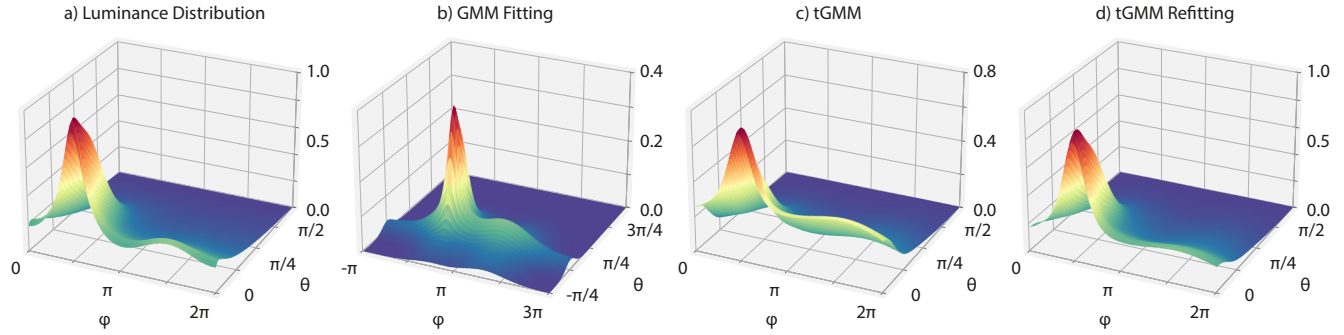

Figure 4: *Our tGMM fitting pipeline. a) Luminance distribution of the sky data to be fitted, b) the resulting GMM from the fitting process; the range of the resulting function exceeds the* $[0,2\pi]\times[0,\frac{\pi}{2}]$ *region, c) the redistribution of the PDF within the truncation bounds and d) the refitting of the weights so that they reflect the new truncated GMM representation.*

Due to its relative high radiance and small subtended solid angle on the hemisphere, most analytic sky models handle the sun's solar disk contribution separately from the main model [HW13, PSS99]. We also perform our fitting on the sky dome illumination without the contribution of the sun disk. Depending on the desired accuracy, the sun is better sampled as a directional light or a spherical ellipsoid [GUnK*17], rather as part of the environment map.

Prior to normalisation and fitting, we weigh the luminance values by $\sin\theta$, the absolute determinant of the Jacobian [BNJ15, PJH16] of the mapping from spherical coordinates (ϕ,θ) to Cartesian coordinates (x,y,z), to account for the non-uniform stretching of the luminance mapping:

$$\hat{L}(u,v)=\frac{L(u,v)}{\max(L)}\sin(\theta). \tag{4}$$

We use the tGMM of Equation 3 as our approximate PDF. We do not follow the formulation of a particular cleat sky model but rather model our PDF as an independent function, with adjustable degrees of freedom that make it directly applicable to any clear sky model. For the model evaluation and practical application of the PDF, we provide fitted tGMM parameters for the Hosek model, which is the state of the art in analytic sky models from ground view, i.e. the most common case. We also provide results for the fitting of the tGMM on the widely used Preetham model.

For the Hosek and Preetham models, we have chosen the number of components N for the tGMM to be 5, after experimentally studying the behaviour of the fitting error (Section 4.2 and Figure 3). Therefore, for each elevation-turbidity parameter pair corresponding to a complete sky map, 25 real values need to be stored for each tGMM. For each of the 5 Gaussians, these values correspond to two μ_k and σ_k parameters, for the longitudinal and latitudinal directions respectively, and one for the component's weight π_k. The storage requirements increase linearly with the required parameter combinations. Since each tGMM requires only the storage of a few real values, it is possible to store large combinations of the radiance-based PDFs of sky model parameters in compact lookup tables.

4.2. Fitting

The PDF fitting operation only needs to be performed once and the resulting tables can be reused across multiple rendering tasks. Therefore, we also provide precomputed tables as part of the supplemental material. We choose to parameterise the GMM directly in the $[0,2\pi]\times[0,\frac{\pi}{2}]$ region instead of normalising it to the $[0,1]^2$ domain. This effectively means that values that are sampled or evaluated through this GMM directly correspond to polar coordinates.

The fitting operation is performed in two steps. The first step uses a robust Trust Region Reflective least squares fitting algorithm [BCL99] to fit a curve defined by a linear superposition of Gaussian kernels on the luminance distribution map (Figure 4a). Initially, we do not constrain our fitted model to be a true GMM, i.e. we only require that weights are positive, but not necessarily summing up to 1. However, the fitted Gaussian components are constrained so that they represent proper normal distributions. Due to the symmetric nature of sky luminance with respect to the parametric space axes, we also enforce a diagonal covariance matrix Σ. This results in fewer parameters and independent sampling of the two random variables θ,ϕ. We further confine the means μ to be inside the region $[0,2\pi]\times[0,\frac{\pi}{2}]$ and variance σ of each component to not exceed 2π. Since we later convert the component distribution to a truncated Gaussian, these constraints are not imperative but rather help avoid arithmetic precision errors during sampling. After the first step completes, the weights are normalised, resulting in a GMM defined on the infinite parametric domain (Figure 4b).

By definition, this GMM is a valid PDF but sampling from it within the specific region of the hemispherical sky dome would require rejection sampling or some unintuitive remapping of the samples. The GMM's domain of support is $\mathbb{R}^2$ and the weights of each component do not correspond to the importance of each Gaussian curve in $[0,2\pi)\times[0,\frac{\pi}{2}]$. We need to convert the GMM into a tGMM effectively making the volume under the curve equal to 1 within the hemispherical parametric space. Therefore, we refit the model to obtain proper weights for truncated Gaussian components. Informally, a truncated normal PDF is constructed by choosing a general normal PDF with μ and σ^2, and a truncation range $[\alpha,\beta]$. The PDF of the associated general normal distribution is modified by setting values outside the range to 0 and uniformly scaling the values in-

Algorithm 1: Sampling from the sky model PDF.

Result: Sampled Direction and PDF
COMPS ← TGMM GAUSSIAN COMPONENTS;
CDF ← 0;
SELECTEDCOMP ← NULL;
$\xi_1, \xi_2, \xi_3 \sim \mathcal{U}(0,1)$;
for $k \leftarrow 1$ **to** N **do**
  CDF += π_k;
  if CDF $\geq \xi_1$ **then**
    SELECTEDCOMP ← COMPS[k];
    break;
  end
end
ϕ=TRUNCGAUSSSAMPLE(SELECTEDCOMP, ξ_2, 0,2π);
θ=TRUNCGAUSSSAMPLE(SELECTEDCOMP, ξ_3, 0, $\frac{\pi}{2}$);
PDF = TRUNCGAUSSPDF(COMPS,ϕ,θ) (Eq. 5)
return TOCARTESIAN($\phi + \phi_{solar}$, θ), PDF.

side the range so that the function integrates to 1. Of course, in the case of a mixture model, the component weights must be refitted, since retaining the initial ones distorts the shape of our PDF (see Figure 4c). To accomplish this, we use a final least squares step to fit the new tGMM. This time, however, we keep μ and σ of the individual Gaussians fixed and only optimise their relative weights. As illustrated in Figure 4d, the resulting shapes of the GMM and the tGMM are very close, as expected; only the weights of the components are different.

The reason we perform the fitting as a two-step process and not as a simultaneous optimisation of all the parameters is the following. With the free variables, namely μ and σ, being present in both the denominator and numerator of the cost function (Equation 2), the function fitting process becomes a much more difficult optimisation problem. Our experiments showed that we could not get as good approximations when fitting directly on a truncated Gaussian mixture model.

We performed the above fitting process on both the Hosek and the Preetham models. For the Hosek model, we generate sky maps using the reference implementation provided by the authors and for the Preetham model we employ a custom implementation. In both cases, the unnormalised radiance distribution is calculated according to Equation 4. As mentioned in Section 4.1, the azimuth of the sun is irrelevant for the PDF parameterisation. However, careful placement facilitates the fitting process. We chose to position the sun at an azimuth of 90 degrees, so that neither the sun halo (90°) nor the opposite brightening of the sky (270°) fall near the sky map range boundaries (see also Figures 2, 4). Otherwise, an additional Gaussian would be required to properly model the illumination transition.

In Figure 6, we present some indicative instances of the absolute error between the normalised luminance (Equation 4) and the fitted tGMM, at various solar elevation and turbidity values. As we can observe, the error rarely reaches 10%, and only at very steep luminance transition zones.

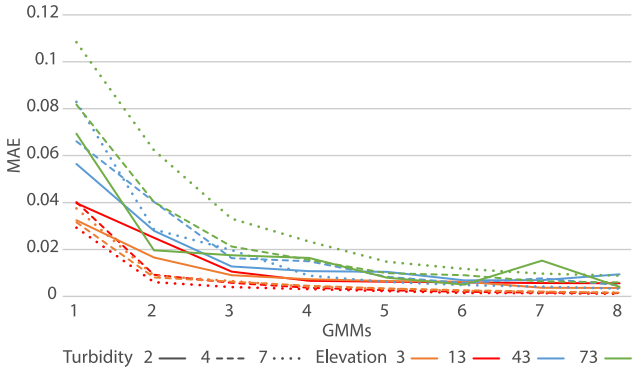

Figure 5: *MAE decreases with the number of Gaussian components, reaching an optimal configuration with 4-6 components.*

Finally, each radiance distribution map is fitted with a selected number of truncated Gaussian components. Our experimental results indicate that 5 Gaussian components adequately model the PDF with a very low approximation error across the parameter space. We present the mean absolute error (MAE) of the fitting with respect to the number of components in Figure 5.

4.3. Sampling and Reconstruction

The process to sample a tGMM is described in Algorithm 1. Using a uniform random variable $\xi_1 \sim U(0,1)$, we construct the 1D CDF of the weights π_k and return the first component where $\xi_1 < \pi_k$. Subsequently, and by exploiting the fact that each component has a diagonal covariance matrix, we can easily sample the 2D truncated Gaussian as two separate 1D truncated Gaussian distributions based on two uniform random variables ξ_2 and ξ_3. For the generation of the random values ξ_1, ξ_2, ξ_3, we follow the sampling process adopted in similar implementations, where two independent variables are generated (ξ_1, ξ_2). One of them is used for the Gaussian lobe selection (ξ_1) and is then remapped to obtain a second uniform random number (ξ_3) for the subsequent (2D) Gaussian sampling.

Given a direction, evaluating the PDF requires the evaluation of the tGMM. Since the model has multiple overlapping regions, the PDF for a specific sample is the weighted sum of the PDF for each component. The PDF needs to also be transformed to be in terms of directions on the unit sphere, which resolves to a simple division operation by the determinant of the Jacobian of the mapping from Cartesian to polar coordinates. The re-weighted values in the stretched polar luminance map and the determinant (approximately) cancel out, leading to perfect importance sampling. The exact calculation of the PDF with respect to the solid angle measure for direction ω comes from Equation 3, where:

$$p(\omega) = \frac{p(\phi,\theta)}{\sin\theta}$$
$$p(\phi,\theta) = \sum_{k}^{N} \pi_\kappa \psi_k(\phi,\theta). \quad (5)$$

One consideration with mixture models is that they may be costly to evaluate if they use a lot of components. In our case, however,

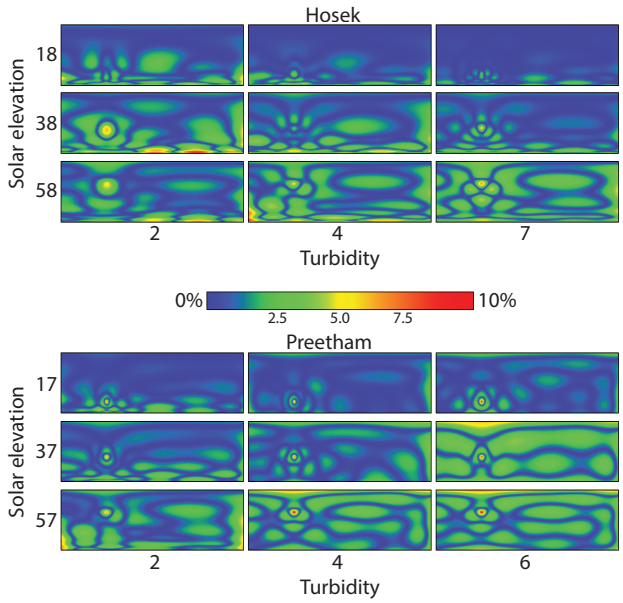

Figure 6: *Absolute error of the fitted model with respect to corresponding normalised sky luminance map for different solar elevation and turbidity values.*

there are ways to further improve performance. The first and most straightforward approach for real-time scenarios and a single sky parameterisation is to precalculate the values and store them in a 2D texture that can be efficiently sampled for each direction. This is a very fast procedure that only needs to be performed when the sky changes and allows for very fast evaluation of the PDF. A more general solution, also suitable for variable sky parameters, is to use a small, precalculated LUT for the standard normal distribution to greatly speed up evaluation time of the exponential functions.

5. Experimental Validation

In our rendering experiments we evaluate our sampling model in terms of both rendering quality, performance and memory requirements. The experimental fitting process verification has been covered in Section 4.2. We compare our method (tGMM) against naïve uniform sampling and the more accurate, but also more computationally expensive SAT-based approach. All experiments in this section use the Hosek model, since it is more physically correct and supports a wider appearance range. All our images were rendered with PBRT-v3, where the additional sampling approaches were integrated in the form of a plugin for infinite light sources.

We provide the full source code of the fitting process as a Python script and precomputed tabulated data for the tGMMs, for both sky models. The Python scripts and the source code for the PBRT plugin can be downloaded from the repository `https://github.com/cgaueb/tgmm_sky_sampling` and constitute a part of the supplemental material. For completeness, a representative set of sky radiance maps, the corresponding PDF maps and rendered images for a variety of scenes and sky model parameters are also provided in the supplemental material.

Quality evaluation. In Figure 7 we show a qualitative comparison among the three sampling approaches on three different scenes, using either Next Event Estimation (NEE) or Multiple Importance Sampling (MIS). Apart from the sky dome, no other light source is present in the scenes. Two of the scenes contain relatively open environments, where direct sky dome lighting is dominant. The *breakfast room* experiment presents a scenario, where the most important (brightest) regions of the sky are directly visible through the opening on the right. We do not employ any form of visibility-based importance sampling here, but rather expect our fitted radiance-based model to accurately place path samples on the important regions during light sampling. We typically use $8/16$ samples (per pixel) for NEE and $4/8$ samples (per pixel/per distribution) in the MIS case. We consider it important to have good results at low sample counts as we are also targeting scenarios with interactive requirements. For each of the resulting images, we include the RMSE and MAE measures. As demonstrated in the rendered examples, our model achieves similar quality to the more expensive SAT approach, but at a significantly lower cost (see performance evaluation next).

As mentioned previously, we perform fitting on a discretisation of the sky model elevation and turbidity. We generate discrete samples of the sky radiance map for different combinations of turbidity and elevation. We sample the elevation every 3 degrees and use the corresponding applicable turbidity ranges per model: 2..10 for the Hosek model and 2..6 for the Preetham model. For arbitrary parameter values between the elevation samples, we assume a piecewise constant behaviour of the PDF, by selecting the tGMM with the closest elevation and turbidity values. We evaluated this choice by snapping the parameters of different skies with already fitted tGMMs to the next closest parameterised tGMM and measuring the difference in produced variance. We observed that even for crude parameter discretisation steps, e.g. 10° in elevation, the resulting variance increase was negligible (below 0.1%).

Performance evaluation. In Tables 1 and 2 we compare the performance of our tGMM model against the SAT-based sampling approach for both CPU and GPU implementations. Note that the SAT implementation requires the full PDF and CDF tables and two binary search queries with an $O(log(n))$ complexity in order to draw a sample. In contrast, our tGMM implementation has complexity $O(N)$, N being the number of Gaussian components (5).

On the CPU, we performed micro-benchmarks that show an estimate of the samples per second one can expect from the reference method implementation in PBRT for SAT-based sampling and our own method. For the SAT-based sampling method we include measurements for both the original map size generated by PBRT (2048×1024) and a scaled-down map size (512×256). Given the nature of GPU devices, it is non-trivial to perform micro-benchmarks of similar accuracy as on the CPU. The reported numbers show the cost of generating 256 samples in a custom CUDA kernel launched for a 1024×512 pixels frame buffer, so that any kernel launch and other computations irrelevant to the benchmark are proportionally amortised. The above timings are for direct illumination only. When the length of the generated paths in-

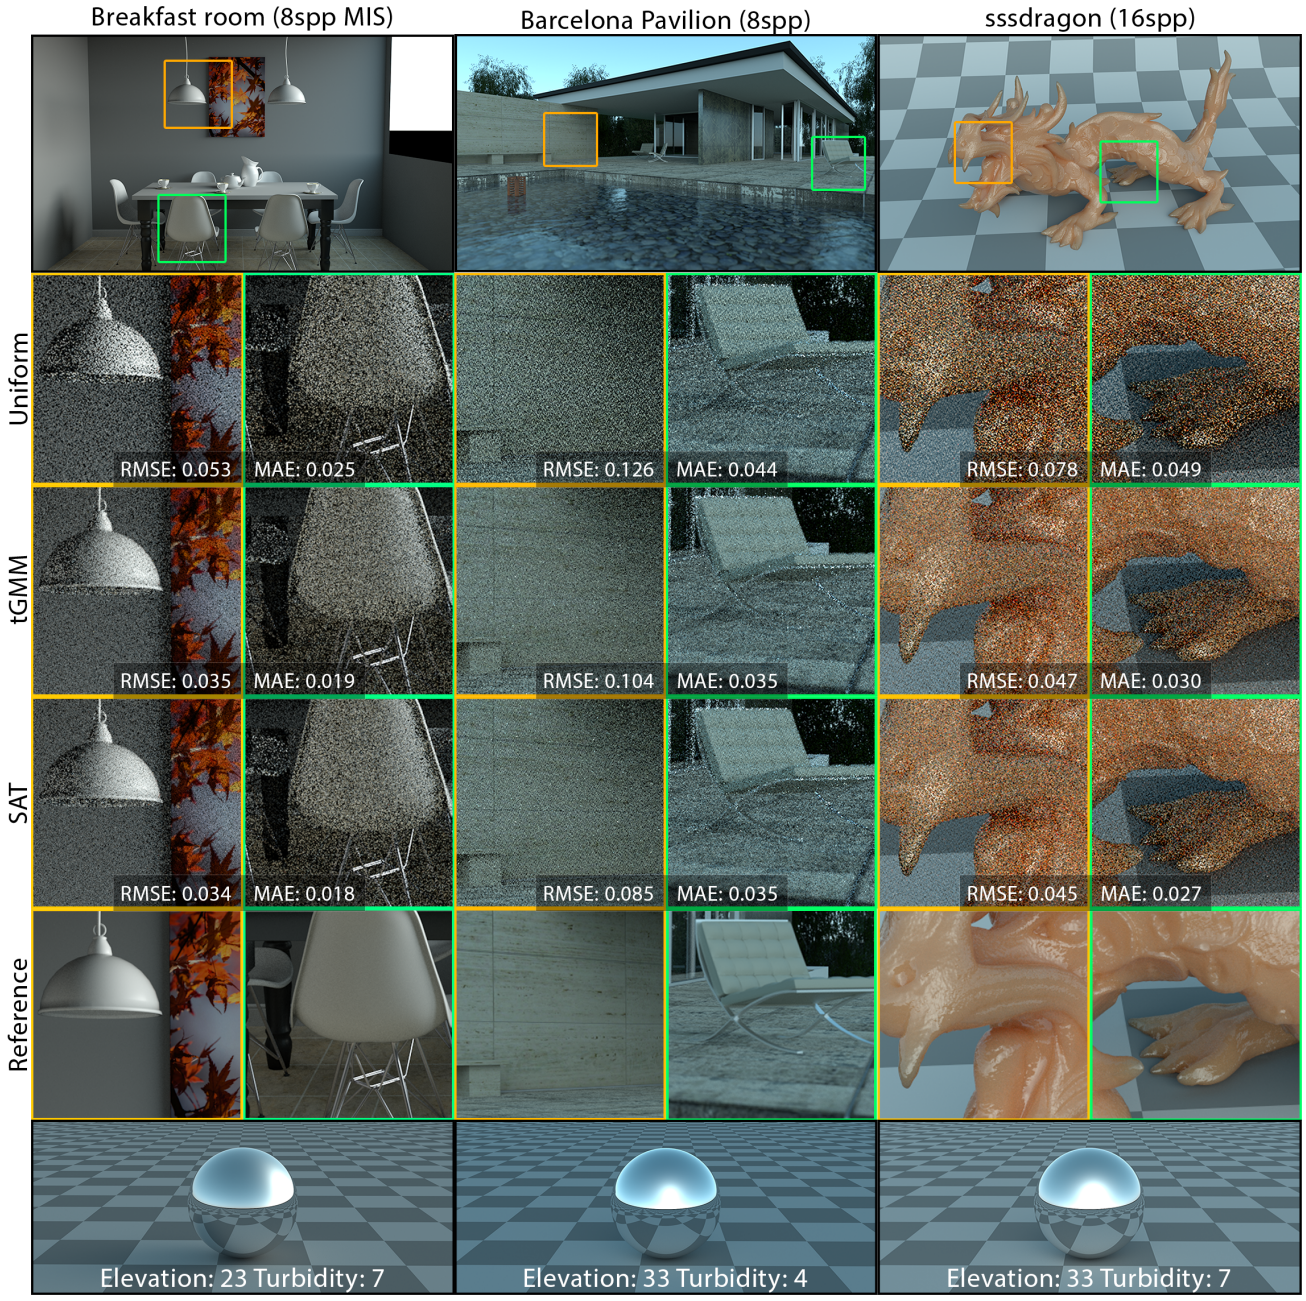

Figure 7: *Qualitative comparison using scenes with different light path complexity, sky model configurations (see images in bottom row) and optional multiple importance sampling enabled, demonstrating the variance reduction of our method relative to naive uniform sampling and the computationally more expensive SAT-based technique.*

creases, the overhead of sampling the sky dome with a SAT-based approach, increases. Measurements are taken from our reference implementation using analytic approximations for the inversion method on Gaussians. Evidently, even for reduced-size SATs, the tGMM model outperforms the SAT-based sampling approach by a large factor and the difference in performance becomes even more noticeable on the GPU implementation.

In Fig. 8 we present equal-quality experiments for the three scenes of Fig. 7. Due to the high precision of our approach, we obtain similar quality to a high-resolution SAT implementation with at most 2 extra samples, resulting in higher overall performance.

Finally, regarding the fitting process, in our current implementation, a minute is required for each sky instance, in the worst case, running on an Intel i7 8700K.

Method	Time (μs)	KSamples/sec	Overhead
SAT 2048	2.868	348.7	2.07X
SAT 512	2.328	429.6	1.68X
tGMM	1.833	545.6	1.33X
Uniform	1.383	721.8	1.00X

Table 1: *Sampling performance comparison on the CPU reported in kilo-samples/sec and time to generate a single sample, averaged over a loop of 1000 iterations.*

Method	Time (ms)	MSamples/sec	Overhead
SAT 2048	263.31	510.3	2.98X
SAT 512	213.33	630.1	2.42X
SAT 128	177.11	758.3	2.01X
tGMM	113.01	1,187.8	1.28X
Uniform	88.14	1,525.2	1.00X

Table 2: *Sampling performance comparison on the GPU reported in giga-samples/sec and measured frame times. Frame times correspond to a pixel shader invocation over a* 1024×512 *frame buffer and 256 sky samples per pixel.*

Memory requirements. For each parameter selection of the model, i.e. turbidity and elevation, only 100 bytes (25 values) are required for nearly optimal luminance-based sampling of the sky. Conversely, the storage requirements for the SAT sampling process on a single sky instance and a typical map resolution of $512 * 256$ is $1MB$ for the conditional PDF and CDF plus $4KB$ for the marginal 1D CDF and PDF. This is a ~$10500\times$ overhead in memory consumption. This overhead can be especially impactful when maintaining and sampling multiple sky instances.

6. Discussion

In this work, we showed that the same properties and intuitions that have led to the formulation of several analytic sky models over the years can be successfully applied to the creation of an equally simple and elegant PDF for sampling such models. We showcased a process that works directly on the radiance distribution data of the sky model to construct a very good approximation using tGMMs. This decouples the fitted PDF from the underlying analytic sky model and can be therefore applied to any, similarly-behaving model, regardless of its formulation and parameterisation. Our experiments showed that this PDF is faster to sample, cheaper to store and offers importance sampling performance very close to that of the more expensive but robust SAT-based alternative.

The focus of our experiments was on clear-sky models, mainly because of their characteristic behaviour. If required, due to the generality of our model and the surface fitting process proposed, it can be fitted with fewer or more Gaussian kernels, effectively modelling different distributions, under diverse storage, run-time performance and quality requirements.

As expected, when the sun elevation measured from the horizon reaches an angle $\theta > 75^\circ$, the clear quality advantage of the model diminishes relative to a simple uniform sampling, because the sky's luminance distribution becomes fairly uniform.

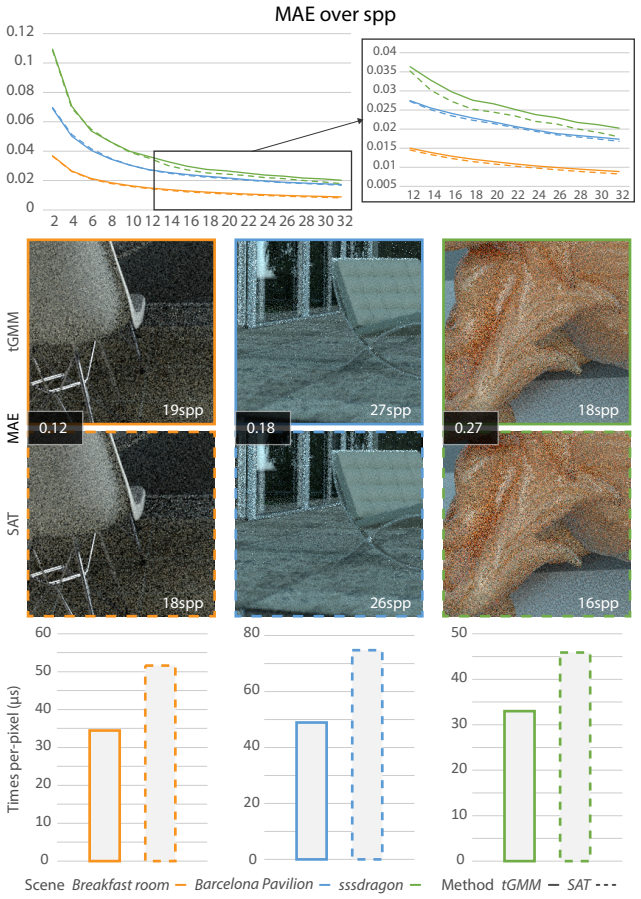

Figure 8: *Equal quality comparison for tGMM vs SAT in the scenes of Figure 7, showing MAE over samples per-pixel (top), equal-quality renders (middle) and CPU timings (bottom).*

In the current method implementation, we fit a tGMM separately for each discrete (elevation, turbidity) parameter pair. As a result, even for similar parameters and respective sky map snapshots, the indices of the Gaussian kernels fitted may randomly switch place from one tGMM to the next, due to the optimisation process. However, certain of the resulting Gaussian components exhibit similar and gradually changing characteristics. For example, one Gaussian component is almost always centred at the circumsolar region and one or two components are "stretched" in order to cover the more uniform regions of the sky or elongated to fit the brightening of the horizon. An interesting next step is to attempt to correlate the Gaussians across time with a similar fitted model that analytically predicts their mean and standard deviation over the parametric space, completely dispensing with tabular data in favour of a fully analytical solution. We would also like to see how this process of tGMM fitting could also be applied to the sky luminance functions themselves.

As a final word, we argue that during this shift from rasterisation to ray tracing, where the importance of proper sampling is significantly elevated, any future attempts at an analytic sky model should

attempt to properly parameterise or define an analytic approach for sampling as well.

Acknowledgements

We would like to thank the anonymous reviewers for their suggestions. The *breakfast room* scene was downloaded from Benedikt Bitterli's rendering resources [Bit16]. All remaining scenes were obtained from the PBRT-v3 scene repository https://www.pbrt.org/scenes-v3.html.

References

[ARBJ03] AGARWAL S., RAMAMOORTHI R., BELONGIE S., JENSEN H. W.: Structured importance sampling of environment maps. In *ACM SIGGRAPH 2003 Papers* (New York, NY, USA, 2003), SIGGRAPH '03, Association for Computing Machinery, p. 605–612. 3

[BCL99] BRANCH M. A., COLEMAN T. F., LI Y.: A subspace, interior, and conjugate gradient method for large-scale bound-constrained minimization problems. *SIAM Journal on Scientific Computing 21*, 1 (1999), 1–23. 5

[Ber11] BERGER M.: Approximate importance sampling of functions reconstructed from spherical harmonics. 3

[Bis06] BISHOP C. M.: *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2006. 2

[Bit16] BITTERLI B.: Rendering resources, 2016. https://benedikt-bitterli.me/resources/. 10

[BNJ15] BITTERLI B., NOVÁK J., JAROSZ W.: Portal-masked environment map sampling. *Comput. Graph. Forum 34*, 4 (July 2015), 13–19. 3, 5

[Bru17] BRUNETON E.: A qualitative and quantitative evaluation of 8 clear sky models. *IEEE Transactions on Visualization and Computer Graphics 23*, 12 (2017), 2641–2655. 3

[C*94] COMMITTEE C., ET AL.: *Spatial distribution of daylight-luminance distributions of various reference skies*. Tech. rep., Technical Report CIE-110-1994, International Commission on Illumination, 1994. 2

[CAM08] CLARBERG P., AKENINE-MÖLLERY T.: Practical product importance sampling for direct illumination. *Computer Graphics Forum 27*, 2 (2008), 681–690. 3

[CETC06] CLINE D., EGBERT P. K., TALBOT J. F., CARDON D. L.: Two stage importance sampling for direct lighting. In *Proceedings of the 17th Eurographics Conference on Rendering Techniques* (Goslar, DEU, 2006), EGSR '06, Eurographics Association, p. 103–113. 3

[CJAMJ05] CLARBERG P., JAROSZ W., AKENINE-MÖLLER T., JENSEN H. W.: Wavelet importance sampling: Efficiently evaluating products of complex functions. *ACM Transactions on Graphics (Proceedings of SIGGRAPH) 24*, 3 (Aug. 2005), 1166–1175. 3

[Cro84] CROW F. C.: Summed-area tables for texture mapping. In *Proceedings of the 11th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1984), SIGGRAPH '84, Association for Computing Machinery, p. 207–212. 3, 4

[GUnK*17] GUILLÉN I., UREÑA C., KING A., FAJARDO M., GEORGIEV I., LÓPEZ-MORENO J., JARABO A.: Area-preserving parameterizations for spherical ellipses. *Comput. Graph. Forum 36*, 4 (July 2017), 179–187. 5

[HEV*16] HERHOLZ S., ELEK O., VORBA J., LENSCH H., KŘIVÁNEK J.: Product importance sampling for light transport path guiding. *Comput. Graph. Forum 35*, 4 (July 2016), 67–77. 3

[Hil20] HILLAIRE S.: A scalable and production ready sky and atmosphere rendering technique. *Computer Graphics Forum 39*, 4 (2020), 13–22. 1, 2

[HMW08] HABEL R., MUSTATA B., WIMMER M.: Efficient Spherical Harmonics Lighting with the Preetham Skylight Model. In *Eurographics 2008 - Short Papers* (2008), Mania K., Reinhard E., (Eds.), The Eurographics Association. 3

[HW12] HOSEK L., WILKIE A.: An analytic model for full spectral sky-dome radiance. *ACM Trans. Graph. 31*, 4 (July 2012). 1, 2

[HW13] HOŠEK L., WILKIE A.: Adding a solar-radiance function to the hošek-wilkie skylight model. *IEEE Computer Graphics and Applications 33*, 3 (2013), 44–52. 5

[JCJ09] JAROSZ W., CARR N. A., JENSEN H. W.: Importance sampling spherical harmonics. *Computer Graphics Forum (Proceedings of Eurographics) 28*, 2 (Apr. 2009), 577–586. 3

[LRR05] LAWRENCE J., RUSINKIEWICZ S., RAMAMOORTHI R.: Adaptive numerical cumulative distribution functions for efficient importance sampling. In *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques* (Goslar, DEU, 2005), EGSR '05, Eurographics Association, p. 11–20. 3

[NDKY96] NISHITA T., DOBASHI Y., KANEDA K., YAMASHITA H.: Display method of the sky color taking into account multiple scattering. In *Pacific Graphics* (1996), vol. 96, pp. 117–132. 1

[NRH03] NG R., RAMAMOORTHI R., HANRAHAN P.: All-frequency shadows using non-linear wavelet lighting approximation. *ACM Trans. Graph. 22*, 3 (July 2003), 376–381. 3

[ODJ04] OSTROMOUKHOV V., DONOHUE C., JODOIN P.-M.: Fast hierarchical importance sampling with blue noise properties. *ACM Trans. Graph. 23*, 3 (Aug. 2004), 488–495. 3

[PJH16] PHARR M., JAKOB W., HUMPHREYS G.: *Physically Based Rendering: From Theory to Implementation*, 3rd ed. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2016. 3, 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. 2

[PSS99] PREETHAM A. J., SHIRLEY P., SMITS B.: A practical analytic model for daylight. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques* (USA, 1999), SIGGRAPH '99, ACM Press/Addison-Wesley Publishing Co., p. 91–100. 1, 2, 5

[RHJD18] REIBOLD F., HANIKA J., JUNG A., DACHSBACHER C.: Selective guided sampling with complete light transport paths. *ACM Trans. Graph. 37*, 6 (Dec. 2018). 3

[SBCD17] SATILMIS P., BASHFORD-ROGERS T., CHALMERS A., DEBATTISTA K.: A machine-learning-driven sky model. *IEEE Computer Graphics and Applications 37*, 1 (2017), 80–91. 3

[SKS02] SLOAN P.-P., KAUTZ J., SNYDER J.: Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Trans. Graph. 21*, 3 (July 2002), 527–536. 3

[TCE05] TALBOT J. F., CLINE D., EGBERT P.: Importance resampling for global illumination. In *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques* (Goslar, DEU, 2005), EGSR '05, Eurographics Association, p. 139–146. 2

[TCJW08] TSAI Y.-T., CHANG C.-C., JIANG Q.-Z., WENG S.-C.: Importance sampling of products from illumination and brdf using spherical radial basis functions. *Vis. Comput. 24*, 7 (July 2008), 817–826. 3

[TS06] TSAI Y.-T., SHIH Z.-C.: All-frequency precomputed radiance transfer using spherical radial basis functions and clustered tensor approximation. *ACM Trans. Graph. 25*, 3 (July 2006), 967–976. 3

[Vea98] VEACH E.: *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford, CA, USA, 1998. AAI9837162. 2

[VKv*14] VORBA J., KARLÍK O., ŠIK M., RITSCHEL T., KŘIVÁNEK J.: On-line learning of parametric mixture models for light transport simulation. *ACM Trans. Graph. 33*, 4 (July 2014). 3

[XWHM20] XIA M. M., WALTER B., HERY C., MARSCHNER S.: Gaussian product sampling for rendering layered materials. *Computer Graphics Forum 39*, 1 (2020), 420–435. 3