
Re-envisioning Sky Models

Ian J. Maquignaz^{†*}, Lucas Valença[†], Aryan Garg[†],
Yannick Hold-Geoffroy[‡], Julien Philip[‡], Jean-François Lalonde[†]
[†]Université Laval, [‡]Adobe Research
*ian.maquignaz.1@ulaval.ca




Figure 1: Laval HDRDB [1]. Mean images were computed from 34,968 images.




Figure 2: Our experimental sky model results.

1 Introduction

A key component of photorealism, natural illumination has a very significant impact on human perception of physical spaces, design of urban architecture, and perceived visual quality of media & film [4]. As a reflection of that importance, the modeling of the sky’s natural illuminations has been a longstanding challenge shared over the past century by astronomers, meteorologist and scientists alike. Early works combined data from varied sources into pre-computed and parametric sky models for engineering and scientific applications, with the first sky models [5, 6] modelling only luminance. With the advent of the digital age, a new paradigm of applications spurred interest in sky models to enable a wide range of emerging digital applications. In this regard, Nishita et al. [7] proposed the first color sky model, enabling the generation of extraterrestrial views of the Earth for space flight simulators, and Image-Based Lighting (IBL) techniques [8] were proposed to render synthetic objects into real and virtual scenes.

2 Background

Though physically captured High Dynamic Range Imagery (HDRI [9]) offers unsurpassed quality and photorealism in the rendering of skies, physical capture is labour-intensive [10] and the resulting imagery is both inflexible and of fixed temporal- and geo-locality. To overcome this rigidity, sky models have progressed greatly over the past four decades with physically-based simulations and parametric models offering a wide range of accurate, versatile, portable and flexible models [11].

To limit the required computational resources, it is often preferable to approximate physically-based simulations of solar and atmospheric illumination via mathematical models and evaluate performance against physically captured skies [10]. One formulation is through the development of numerical models such as [7, 12, 13, 14, 15, 16], which derive simplified mathematical representations for complex atmospheric systems. Though reducing computation expense, numerical models generally remain complex, memory intensive, and require a pre-computation step. An alternative formulation introduced by Perez et al. [6] is the fitting of an analytical model to a body of sky data. These simpler models can be fitted to computationally expensive, accurate, and diverse data acquired from complex models (e.g., Preetham et al. [17] is fitted to [12]), [18, 19, 20] use path tracers such as A.R.T [21], and Bruneton [11] evaluated against libRadtran [22] physical simulations and Kider [10] physically captured dataset. Such models trade-off accuracy to produce lightweight, flexible, and fast parametric models which support a wide range of applications [11].

With the advent of deep learning, the fitting of analytical models to features and modalities can be automated, with parameterization to a finite set of latent parameters, or to intuitive user-parameters by supervised learning. Notably, this concept has been proposed for lighting estimation [23, 24, 25, 26, 27], where LDR images can be used to guide the generation of HDR environment maps for relighting virtual objects and scenes. Though these learned representations tend to offer little photorealism and generate overly smooth skies, they capture lighting energy with significant accuracy. Most recently, approaches have been proposed to augment parametric skies with photorealistic clouds. Satilmis et al. [28] proposed a method for augmenting Hošek-Wilkie [18] parametric skies per user-controlled cloud placement. SkyGAN [29] proposed fitting the Hošek-Wilkie [20] model to real-world photographs, enabling the re-generation of Hošek-Wilkie skies with atmospheric formations, though with limited dynamic range and photorealism in atmospheric formations.

Despite this progress, sky models are predominantly limited to the illumination of clear, hazy, and/or overcast daylight skies and are unable to approximate non-uniform atmospheric formations such as clouds. Though parametric models allow for flexibility and versatility through a finite set of user-parameters, additional steps are required to include complex atmospheric formations, with proposed steps including volumetric cloud rendering [30] and cloud simulations [31]. While this combined approach can be versatile and photorealistic, it can be labour intensive, computationally expensive and difficult to configure [32]. This approach is satisfactory for many applications including generating skies for plausible extra-solar worlds but, for realism and weather variations, the physically captured HDRI remains unsurpassed [33, 34].

3 Methodology

In this document, we present a novel clear sky model (SkyNet) together with recent results for user-controlled generation of atmospheric formations. We further showcase learned sky models capable of producing photorealistic weathered skies with faithful and controllable HDR illumination, learned directly from physical sky captures. Our methods combine the versatility of parametric models and the realism of deep generative networks.

The Laval HDR Sky database (HDRDB, [1]) consists of 34K+ HDR images captured in Quebec City, Canada across varied time intervals between 2014 and 2016 using the capture method proposed by Stumpf et al. [35]. From the database—the mean of which is shown in Fig. 1c—we augment the dataset with random rotations around the zenith to increase solar placement coverage (Fig. 1d. See C.1 for nomenclature) and enable generation of skies outside of HDRDB. To handle the extreme HDR values present in the database, we apply logarithmic tone-map using $I' = \log_2(I + 1)$ (see C.3.1 for details). Solar masks were created from ephemeris calculations [36] and cloud masks were thresholded from the color ratio $Y = \frac{B-R}{B+R}$ proposed by Dev et al. [37].

We have trained four different GAN pipelines on HDRDB: SPADE [2], SEAN [3], and our clear and weathered models (both UNetFixup [38]). In SPADE and SEAN, we observe some undesired clouds are present in the output due to the models learning positional artifacts from the dataset. We address this in our clear SkyNet model by implementing an L1 loss which guides the model toward a median cloudless image with no visible artifacts. This can be seen in SkyNet’s generation of Fig. 2a per a label created from Fig. 1a). To create the weathered model, we leverage the Lalonde-Matthews (LM) outdoor illumination model [39] to parameterize the UNetFixup network to 11 intuitive parameters.

We observe that SPADE (Fig. 2c) collapses to a singular modality of atmospheric formations, while SEAN (Fig. 2e) and our weathered model produce novel, uniform, and visually-appealing skies.

4 Conclusion

Through our showcased works, we demonstrate sky modeling with Deep Neural Networks (DNNs), leveraging physically captured HDRI to learn physically accurate skies with diverse and photorealistic atmospheric formations. We demonstrate that our models, both clear and weathered, generate global representations of HDR skies per user-controlled label maps, from sunrise to sunset. We achieve state-of-the-art diversity, fidelity and photorealism, and offers versatility in the emulation of locations, dates, times and various weather conditions. Where previous works have faced trade-offs between with atmospheric formations and computational cost, we demonstrate encompassing models which generates weathered HDRI ready for 3D rendering engines.

References

- [1] J.-F. Lalonde, L.-P. Asselin, J. Becirovski, Y. Hold-Geoffroy, M. Garon, M.-A. Gardner, and J. Zhang. (2016) The Laval HDR sky database. [Online]. Available: <http://sky.hdrdb.com>
- [2] T. Park, M. Liu, T. Wang, and J. Zhu, “Semantic image synthesis with spatially-adaptive normalization,” *CoRR*, vol. abs/1903.07291, 2019. [Online]. Available: <http://arxiv.org/abs/1903.07291>
- [3] P. Zhu, R. Abdal, Y. Qin, and P. Wonka, “Sean: Image synthesis with semantic region-adaptive normalization,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [4] K. Alshaibani and D. Li, “Sky type classification for the ISO/CIE Standard General Skies: a proposal for a new approach,” *International Journal of Low-Carbon Technologies*, vol. 16, no. 3, pp. 921–926, 03 2021. [Online]. Available: <https://doi.org/10.1093/ijlct/ctab020>
- [5] P. Moon, “Proposed standard solar-radiation curves for engineering use,” *Journal of the Franklin Institute*, vol. 230, no. 5, pp. 583–617, 1940. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0016003240903647>
- [6] R. Perez, R. Seals, and J. Michalsky, “All-weather model for sky luminance distribution—preliminary configuration and validation,” *Solar Energy*, vol. 50, no. 3, pp. 235–245, 1993. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0038092X93900171>
- [7] T. Nishita, T. Sirai, K. Tadamura, and E. Nakamae, “Display of the earth taking into account atmospheric scattering,” in *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, 1993, pp. 175–182.
- [8] P. Debevec, “Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography,” in *Proc. SIGGRAPH 98*, 1998, pp. 189–198.
- [9] E. Reinhard, W. Heidrich, P. Debevec, S. Pattanaik, G. Ward, and K. Myszkowski, *High dynamic range imaging: acquisition, display, and image-based lighting*. Morgan Kaufmann, 2010.
- [10] J. T. Kider, D. Knowlton, J. Newlin, Y. K. Li, and D. P. Greenberg, “A framework for the experimental comparison of solar and skydome illumination,” *ACM Trans. Graph.*, vol. 33, no. 6, nov 2014. [Online]. Available: <https://doi.org/10.1145/2661229.2661259>
- [11] E. Bruneton, “A qualitative and quantitative evaluation of 8 clear sky models,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 12, pp. 2641–2655, 2017.
- [12] T. Nishita, Y. Dobashi, and E. Nakamae, “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*, 1996, pp. 379–386.
- [13] S. O’Neil, “Accurate atmospheric scattering,” *Gpu Gems*, vol. 2, pp. 253–268, 2005.
- [14] J. Haber, M. Magnor, and H.-P. Seidel, “Physically-based simulation of twilight phenomena,” *ACM Transactions on Graphics (TOG)*, vol. 24, no. 4, pp. 1353–1373, 2005.
- [15] E. Bruneton and F. Neyret, “Precomputed atmospheric scattering,” *Computer Graphics Forum*, vol. 27, no. 4, pp. 1079–1086, 2008.
- [16] O. Elek and P. Knoch, “Real-time spectral scattering in large-scale natural participating media,” in *Proceedings of the 26th Spring Conference on Computer Graphics*, 2010, pp. 77–84.
- [17] A. J. Preetham, P. Shirley, and B. Smits, “A practical analytic model for daylight,” in *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH ’99. USA: ACM Press/Addison-Wesley Publishing Co., 1999, p. 91–100. [Online]. Available: <https://doi.org/10.1145/311535.311545>
- [18] L. Hosek and A. Wilkie, “An analytic model for full spectral sky-dome radiance,” *ACM Trans. Graph.*, vol. 31, no. 4, jul 2012. [Online]. Available: <https://doi.org/10.1145/2185520.2185591>
- [19] L. Hošek and A. Wilkie, “Adding a solar-radiance function to the hošek-wilkie skylight model,” *IEEE Computer Graphics and Applications*, vol. 33, no. 3, pp. 44–52, 2013.

- [20] A. Wilkie, P. Vevoda, T. Bashford-Rogers, L. Hošek, T. Iser, M. Kolářová, T. Rittig, and J. Křivánek, “A fitted radiance and attenuation model for realistic atmospheres,” *ACM Trans. Graph.*, vol. 40, no. 4, Jul. 2021. [Online]. Available: <https://doi.org/10.1145/3450626.3459758>
- [21] The advanced rendering toolkit. [Online]. Available: https://cgg.mff.cuni.cz/ART/archivers/art_2_1_1.html
- [22] C. Emde, R. Buras-Schnell, A. Kylling, B. Mayer, J. Gasteiger, U. Hamann, J. Kylling, B. Richter, C. Pause, T. Dowling, and L. Bugliaro, “The libradtran software package for radiative transfer calculations (version 2.0.1),” *Geoscientific Model Development*, vol. 9, no. 5, pp. 1647–1672, 2016. [Online]. Available: <https://gmd.copernicus.org/articles/9/1647/2016/>
- [23] Y. Hold-Geoffroy, K. Sunkavalli, S. Hadap, E. Gambaretto, and J.-F. Lalonde, “Deep outdoor illumination estimation,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 7312–7321.
- [24] Y. Hold-Geoffroy, A. Athawale, and J.-F. Lalonde, “Deep sky modeling for single image outdoor lighting estimation,” in *CVPR*, 2019.
- [25] J. Zhang, K. Sunkavalli, Y. Hold-Geoffroy, S. Hadap, J. Eisenman, and J.-F. Lalonde, “All-weather deep outdoor lighting estimation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10 158–10 166.
- [26] P. Yu, J. Guo, F. Huang, C. Zhou, H. Che, X. Ling, and Y. Guo, “Hierarchical disentangled representation learning for outdoor illumination estimation and editing,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15 313–15 322.
- [27] P. Yu, J. Guo, L. Wu, C. Zhou, M. Li, C. Wang, and Y. Guo, “Dual attention autoencoder for all-weather outdoor lighting estimation,” *Sci. China Inf. Sci.*, vol. 64, 2021.
- [28] P. Satilmis, D. Marnierides, K. Debattista, and T. Bashford-Rogers, “Deep synthesis of cloud lighting,” *IEEE Computer Graphics and Applications*, 2022.
- [29] M. Mirbauer, T. Rittig, T. Iser, J. Křivánek, and E. Šikudová, “Skygan: Towards realistic cloud imagery for image based lighting,” in *Eurographics Symposium on Rendering. The Eurographics Association*, 2022.
- [30] S. Kallweit, T. Müller, B. McWilliams, M. Gross, and J. Novák, “Deep scattering: Rendering atmospheric clouds with radiance-predicting neural networks,” *ACM Trans. Graph. (Proc. of Siggraph Asia)*, vol. 36, no. 6, Nov. 2017.
- [31] A. Bouthors, F. Neyret, N. Max, E. Bruneton, and C. Crassin, “Interactive multiple anisotropic scattering in clouds,” in *Proceedings of the 2008 symposium on Interactive 3D graphics and games*, 2008, pp. 173–182.
- [32] S. Hill, S. McAuley, L. Belcour, W. Earl, N. Harrysson, S. Hillaire, N. Hoffman, L. Kerley, J. Patry, R. Pieké, I. Skliar, J. Stone, P. Barla, M. Bati, and I. Georgiev, “Physically based shading in theory and practice,” in *ACM SIGGRAPH 2020 Courses*, 2020.
- [33] A. Wilkie and L. Hošek, “Predicting sky dome appearance on earth-like extrasolar worlds,” in *Proceedings of the 29th Spring Conference on Computer Graphics*, ser. SCCG ’13. New York, NY, USA: Association for Computing Machinery, 2013, p. 145–152. [Online]. Available: <https://doi.org/10.1145/2508244.2508263>
- [34] Forza horizon 5’s art team explains how it got the skies to look so realistic. [Online]. Available: <https://www.videogameschronicle.com/news/forza-horizon-5s-art-team-explains-how-it-got-the-skies-to-look-so-realistic/>
- [35] J. Stumpfel, A. Jones, A. Wenger, C. Tchou, T. Hawkins, and P. Debevec, “Direct hdr capture of the sun and sky,” in *ACM SIGGRAPH 2006 Courses*, ser. SIGGRAPH ’06. New York, NY, USA: Association for Computing Machinery, 2006, p. 5–es. [Online]. Available: <https://doi.org/10.1145/1185657.1185687>
- [36] Pysolar. [Online]. Available: <https://github.com/pingswept/pysolar>
- [37] S. Dev, F. M. Savoy, Y. H. Lee, and S. Winkler, “High-dynamic-range imaging for cloud segmentation,” *Atmospheric Measurement Techniques*, vol. 11, no. 4, pp. 2041–2049, 2018. [Online]. Available: <https://amt.copernicus.org/articles/11/2041/2018/>
- [38] D. Griffiths, T. Ritschel, and J. Philip, “Outcast: Outdoor single-image relighting with cast shadows,” *Computer Graphics Forum*, vol. 41, no. 2, pp. 179–193, 2022. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14467>

- [39] J.-F. Lalonde and I. Matthews, "Lighting estimation in outdoor image collections," in *2014 2nd International Conference on 3D Vision*, vol. 1. IEEE, 2014, pp. 131–138.
- [40] Blender Foundationn. Blender 3.1. [Online]. Available: <https://www.blender.org/>

A Abstract (english):

Sky models are an integral part of environment simulations with applications in civil engineering, urban planning, and visual arts. Recent works have extended sky models to be more comprehensive and, with the advent of deep learning, interest has emerged in evolving towards Generative Adversarial Networks. This work proposes a novel deep learning approach to sky models, generating skies per user-controlled positioning of solar and atmospheric components, for photorealistic emulation of locations, dates, and weather conditions.

B Abtrait (français):

La modélisation du ciel est une partie intégrante des simulations d'environnement, avec des applications en génie civil, en urbanisme et en arts visuels. Des travaux récents ont étendu les modèles du ciel pour qu'ils soient plus complets et, avec l'avènement de l'apprentissage en profondeur, l'évolution de ces modèles vers les réseaux antagonistes génératifs. Ce travail propose une nouvelle approche d'apprentissage en profondeur pour les modèles du ciel, générant des ciels avec positionnement des composants solaires et atmosphériques contrôlé par un utilisateur, pour une émulation photoréaliste des lieux, des dates et des conditions météorologiques variées.

C Appendix

C.1 Skydome Nomenclature

As illustrated in Figure 1, Laval HDRDB consists of images captured with a calibrated fisheye camera. As part of calibration for camera intrinsics and extrinsics, the images are transformed to latlong environment maps as shown in Figure 3c. This format is sub-optimal for DNN modelling, as the $360^\circ N \times 2N$ panoramic representation introduces a large zeroed region below the horizon and fails to convey continuity (skies are split by three borders). To accommodate, the latlong environment maps are transformed to sky-angular format as defined by the solid angles in Figure 3b and shown in Figure 3d.

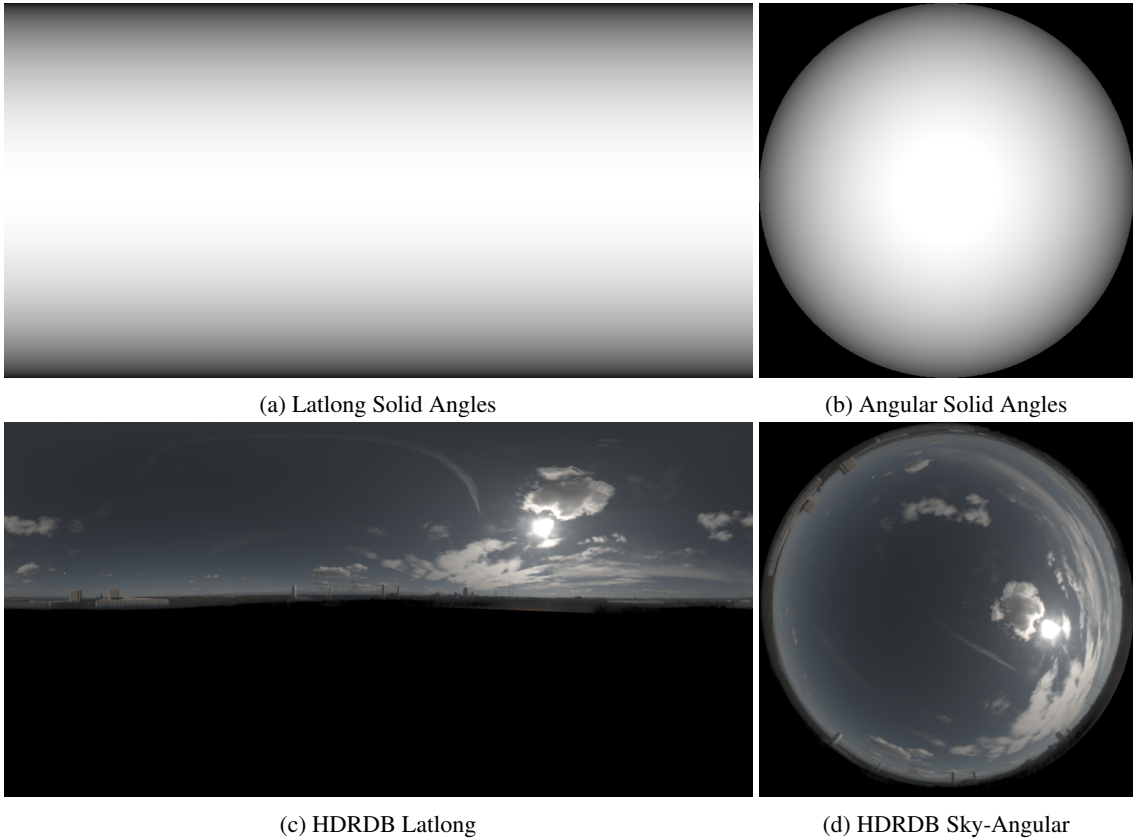


Figure 3: Latlong ($N \times 2N$) and angular ($N \times N$) environment maps

Sky-angular skydomes are assumed perfectly hemispherical with the zenith centred on the image plane. This allows for the definition of solar positioning in terms of solar elevation (φ) and azimuth angle (θ) as illustrated in Figure 4.

The image plane coordinate system denoted in terms of $+X$ columns and $+Y$ rows of pixels in Figure 4, illustrates the inequality in the skydome ‘surface area’ sampled by each pixel. As illustrated in Figures 3a and 3b, this inequality is reflected in a non-linear distribution of radiant energy which can be quantized in terms of solid angles. In transformation and resizing of latlong and sky-angular skydomes, the solid angles must be accounted for in order retain the cumulative radiant energy of solar and diffuse light.

In our work, we first inter-area down-sample the $N \times 2N$ latlong skydomes by a factor of 2^x before transforming to sky-angular representation. In experimentation this approach provided the best retention of radiant energy, with the latlong environment map shown in Figure 3c retaining 99.9988% of its energy after down-sampling from 1024×2048 to 512×1024 , and in turn retaining 99.79% of the original energy after transformation to 512×512 sky-angular representation as shown in Figure 3d.

C.2 Parametric Clear Skies

Figure 5 illustrates three parametric clear skies rendered using the Blender Creation Suite [40] with skies generated as scene backgrounds for fisheye lens ‘capture’. Where possible, all clear sky parameters were kept constant.

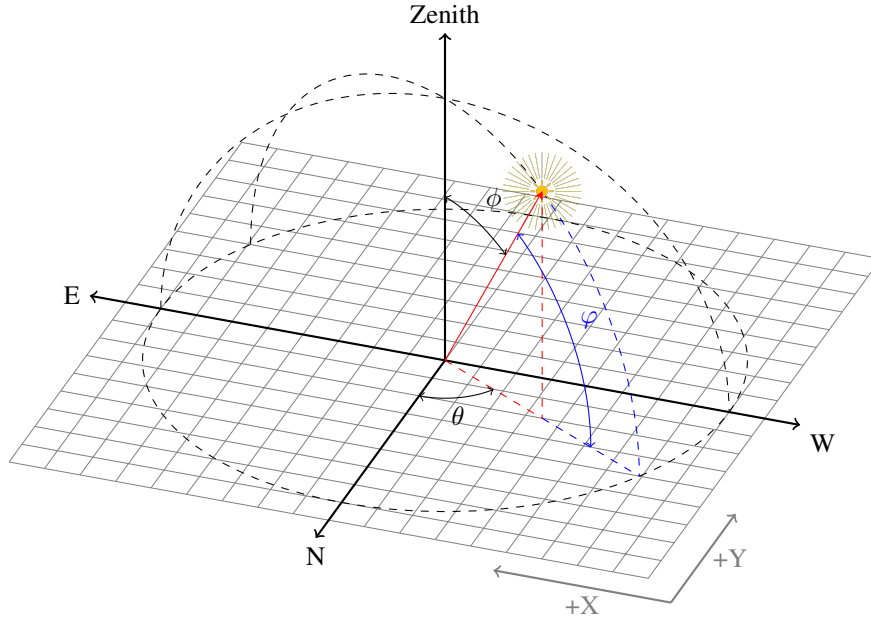
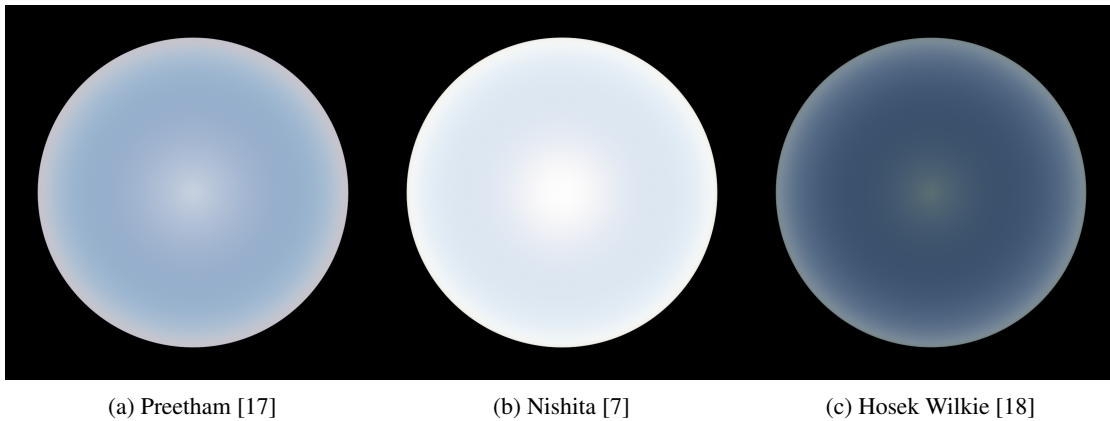


Figure 4: Definition of a skyangular skydome for solar angle (φ) and Azimuth angle (θ)



(a) Preetham [17]

(b) Nishita [7]

(c) Hosek Wilkie [18]

Figure 5: Parametric Clear Skies rendered in Blender [40]

C.3 Validation of Dynamic Range

C.3.1 Tone-Mapping

Figure 6 illustrates the dynamic range exhibited by clear skies (Figure 6a), and the impact of tone-mapping (Figure 6b) in luminance space Y . We compress the dynamic range of each HDRI I_{in} using global logarithmic tone-mapping $I' = \log_2(I + 1)$. This effectively reduces the dynamic range by an approximate factor of $\times 10,000$, which can be inverted via $I = 2^{(I')} - 1$. Our models are trained directly in this logarithmic space, and we apply the inverse tone-mapping on the output to obtain an HDRI that is ready to use by 3D rendering engines.

During training, we do not clip the dynamic range of the image I' to the range $[0, 1]$ to mitigate loss in the reconstructed dynamic range of the image I . For visualization, images illustrated in this paper are Gamma 2.2 tone-mapped and clipped to the range $[0, 1]$.

To the best of our abilities we have balanced the subsets, discrepancies exist due to solar elevation at time of capture. Though not illustrated herien, histograms have also been created to determine subset balance given solar elevation, solar azimuth, and cloud cover. As shown in Figure 6a, the test subset of the data has a greater portion of data captured during summer months (higher solar elevation) and is therefore cumulatively brighter. Spikes in the train and test subsets are believed to be the result of rain events.

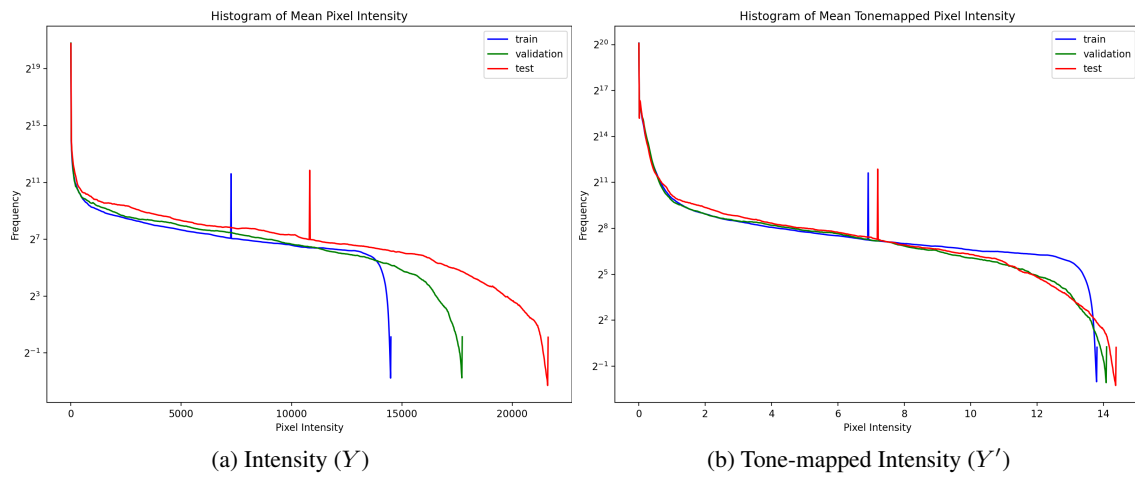


Figure 6: Histograms of HDR Skydomes from the Laval HDRDB illustrating the number of pixels (y-axis) against their intensity (x-axis) for our training, test and validation subsets. Figure 6b illustrates the impact of tone-mapping, without clipping of the tone-mapped range