

DeepSky

Ian J. Maquignaz¹ Aryan Garg^{1,2} Jean-François Lalonde¹

¹Département de Génie Électrique, Université Laval, Québec, QC G1V 0A6

²MITACS Globalink Research Program, Canada

Sky models are an integral part of daylight environment simulation with various applications, including civil engineering, urban planning and visual arts. Recent works have extended parametric clear sky models to be more comprehensive and, with the emergence of deep learning, interest has emerged in evolving these models with Generative Adversarial Networks. This work proposes a novel deep learning approach to sky models, enabling generation of skies per user-controlled positioning of solar and atmospheric components, for visually appealing emulation of locations, dates, times and various weather conditions.



(a) Clear Sky (b) Cloudy Sky (c) Mean Sky (d) Aug. Mean Sky (e) SkyNet (ours)

Figure 1: Laval HDRDB and SkyNet (ours). Mean images were computed from 34,968 images.

1 Introduction

Modelling the sky has been a longstanding challenge shared over the past century by astronomers, meteorologist and scientists alike. Early works aimed to combine data from these varied sources into usable pre-computed and parametric sky models, with one of the first proposed models by Moon and later adopted as part of a ‘standard overcast sky’ by the *International Commission on Illumination* (CIE¹) in 1955 [1]. Intermediary models were proposed to fill gaps between overcast and clear skies, with the first “all weather” illumination model proposed by Perez et al. in 1993 [2], and the first color model proposed by Nishita in 1993 with the goal of enabling space flight simulators [3]. In recent years, parametric models have grown in accuracy and reduced computational time requirements (see A.4)[4].

Despite this progress, existing sky models remain approximations aimed at modelling the atmosphere’s clear and/or overcast daylight illumination of the earth’s surface. With variations in particle and aerosol compositions, state-dependent thermodynamic properties which define scattering, and surface albedo (e.g. urban, vegetation, and maritime environments), quantifying the sky is beyond the collective capacities of today’s instruments [5]. Frameworks have been proposed for sampling solar and diffuse illumination, but collection is difficult and remains accurate only for the temporal and geographical locality in which it was captured [6]. To provide accuracy and variability, physical simulations of solar and thermal radiation in the earth’s atmosphere have been created in the libRadtran [7] and A.R.T [8] frameworks. Though proven precise, these simulations are computationally expensive, making approximation with pre-computed and parametric models preferable [4]. This approach is satisfactory for many applications including generating skies for plausible extra-solar worlds but, for realism and weather variations, the real-world captured by HDRI remains unsurpassed [9, 10].

¹Commission Internationale de l’Éclairage

This work proposes a novel deep neural network (DNN) approach to sky models, leveraging the Laval HDR Sky database and tailoring it to a model which is readily compatible with segmentation maps of atmospheric components. We demonstrate that our model, SkyNet, generates physically accurate and visually appealing global representations of HDR skies per user-controlled solar label maps. This approach achieves higher fidelity and resolution than previous works, and offers versatility in the emulation of locations, dates, times and various weather conditions. Where previous works have modelled clear and overcast skies, our work paves the way for the creation of the first partially cloudy sky model.

2 Method

We propose a clear sky model with controlled atmospheric components learned from HDRI. **Architecture.** We implement the OutCast outdoor image relighting model proposed by Griffiths et al. [11]. The input to our model is a label mask with the sun location encoded as a cosine distance map $M_d \in \mathbb{R} \cap [-1, 1]$, an example of which is shown in Fig. 2a. Our model directly outputs a skydome image ready to be employed as HDRI in a renderer. We train our model using an $L1$ loss and the Adam optimizer [12] with an initial learning rate of 1×10^{-4} . Once trained, the model requires only the input of a solar label mask, enabling easy validation with user-controlled input.

Data. To train our model, we employ the Laval HDR Sky database (Laval HDRDB, [13]). This dataset consists of 34K+ HDR images captured in Quebec City, Canada across varied time intervals between 2015 and 2016 using the capture method proposed by Stumpfel et al. [14]. HDRDB consists of physically accurate and calibrated linear environment maps, which we convert to the sky-angular format (see A.3). As illustrated in Fig. 1, we manually split the database into clear and cloudy skies. To prevent a temporal imbalance, many skies labelled as clear still bear some sparse clouds, however, our $L1$ loss is able to ignore those clouds by guiding the model towards a median cloud-less image [15, 16]. 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) and enable generation of skies outside of HDRDB (Fig. 1e). To handle the extreme HDR values present in the database, we apply logarithmic tone-map using $I_{\text{out}} = \log_2(I_{\text{in}} + 1)$ (see A.5.1 for details).

Results. When training an existing generative imaging architecture such as Pix2PixHD [17] on our data, we observe some undesired clouds are present in the output (Fig. 2b). In contrast, our method produces much more uniform and visually appealing clear skies (Fig. 2c and Fig. 3). We validate that our model preserves the physical accuracy of our real-sky dataset with metrics including Earth Mover Distance (EMD), χ^2 , and Kullback Leibler divergence, the results of which are shown in A.5.2.




Figure 2: Generated clear skies

3 Discussion

In this work, we propose a physically accurate sky model directly learned from captured data. Our model allows users to control the sun position and produces visually pleasing clear skies. We develop our model using a global representation, the cosine distance map, on which cloud maps can be overlaid in a straightforward manner. This makes our model a first step towards generating photorealistic cloudy skies. In comparison, existing DNN models are much less adapted to this future goal: the clear sky MLP from Satilmis et al.




Figure 3: SkyNet (ours) can produce physically accurate clear skies

[18] uses a local representation and will not consider the spatial interactions generated by clouds. Similarly, the outputs of Hold-Geoffroy et al. [19] are hindered by low resolution and a different sky representation that amplifies spherical distortions above the camera, where clouds are likely to be.

References

- [1] 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>
- [2] 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/0038092X9390017I>
- [3] 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.
- [4] E. Bruneton, “A qualitative and quantitative evaluation of 8 clear sky models,” *CoRR*, vol. abs/1612.04336, 2016. [Online]. Available: <http://arxiv.org/abs/1612.04336>
- [5] M. Inanici, B. Abboushi, and S. Safranek, “Evaluation of sky spectra and sky models in daylighting simulations,” *Lighting Research & Technology*, vol. 0, no. 0, p. 14771535221103400, 0. [Online]. Available: <https://doi.org/10.1177/14771535221103400>
- [6] 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>
- [7] 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/>
- [8] The advanced rendering toolkit. [Online]. Available: https://cgg.mff.cuni.cz/ART/archivers/art_2_1_1.html
- [9] 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>
- [10] 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/>
- [11] 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>
- [12] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [13] 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>

- [14] 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>
- [15] J. Lehtinen, J. Munkberg, J. Hasselgren, S. Laine, T. Karras, M. Aittala, and T. Aila, “Noise2noise: Learning image restoration without clean data,” *arXiv preprint arXiv:1803.04189*, 2018.
- [16] D. Ulyanov, A. Vedaldi, and V. Lempitsky, “Deep image prior,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 9446–9454.
- [17] T. Wang, M. Liu, J. Zhu, A. Tao, J. Kautz, and B. Catanzaro, “High-resolution image synthesis and semantic manipulation with conditional gans,” *CoRR*, vol. abs/1711.11585, 2017. [Online]. Available: <http://arxiv.org/abs/1711.11585>
- [18] P. Satilmis, T. Bashford-Rogers, K. Debattista, and A. Chalmers, “A machine learning driven sky model,” *IEEE Computer Graphics and Applications*, vol. 37, pp. 1–1, 05 2016.
- [19] Y. Hold-Geoffroy, A. Athawale, and J.-F. Lalonde, “Deep sky modeling for single image outdoor lighting estimation,” in *CVPR*, 2019.
- [20] 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.
- [21] H. W. Jensen, F. Durand, J. Dorsey, M. M. Stark, P. Shirley, and S. Premože, “A physically-based night sky model,” in *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH ’01. New York, NY, USA: Association for Computing Machinery, 2001, p. 399–408. [Online]. Available: <https://doi.org/10.1145/383259.383306>
- [22] H.-Y. Zhou, B.-B. Gao, and J. Wu, “Sunrise or sunset: Selective comparison learning for subtle attribute recognition,” 2017. [Online]. Available: <https://arxiv.org/abs/1707.06335>
- [23] Blender Foundationn. Blender 3.1. [Online]. Available: <https://www.blender.org/>
- [24] 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>
- [25] 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>
- [26] G. Bradski, “The OpenCV Library,” *Dr. Dobb’s Journal of Software Tools*, 2000.
- [27] J. Puzicha, T. Hofmann, and J. Buhmann, “Non-parametric similarity measures for unsupervised texture segmentation and image retrieval,” 07 1997, pp. 267–272.
- [28] A. Ramdas, N. Garcia, and M. Cuturi, “On wasserstein two sample testing and related families of nonparametric tests,” 2015. [Online]. Available: <https://arxiv.org/abs/1509.02237>

- [29] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods*, vol. 17, pp. 261–272, 2020.

A Appendix

A.1 3-Sentence Abstract

Sky models are an integral part of daylight environment simulation with various applications, including civil engineering, urban planning and visual arts. Recent works have extended parametric clear sky models to be more comprehensive and, with the emergence of deep learning, interest has emerged in evolving these models with Generative Adversarial Networks. This work proposes a novel deep neural network (DNN) approach to sky models, enabling generation of HDR skies per user-controlled positioning of solar and atmospheric components in label maps, creating visually appealing emulation of locations, dates, times and various weather conditions.

A.2 Further Discussion

Though our proposed model allows for user-controlled solar placement and thus the simulation of skies for any geolocation, the inherent atmospheric effects (i.e. those specified through user-controlled label maps and learned by our model) are learned from the Laval HDRDB and thus not necessarily realistic for localities other than Quebec City.

Though offering near complete coverage, our proposed model does not permit for the placement of the sun near the zenith due to a lack of samples in said region. This is illustrated in Figure 1 by following the solar paths (lines) in Figure 1c across the south, resulting in a gap around the zenith in Figure 1d. This erroneous region for solar placement is the direct result of Quebec City being geolocated in the northern hemisphere, and thus not on the equator.

High Dynamic Range Imagery (HDRI) requires special consideration to capture the 18 f-stops of irradiance without saturation (over- or under-exposure) [14, 20]. To our knowledge, no database of HDRI skydomes encompasses geographical variation and/or extends beyond capture of daylight (e.g. lunar, starlight, and etc [21]).




Figure 4: SkyNet paired clear skies from June 25th, 2015. Top row is cosine distance solar masks, middle row is images from HDRDB and bottom row is the corresponding clear skies generated by SkyNet.




Figure 5: SkyNet unpaired clear skies with solar positioning via cosine distance masks

Though the dataset was manually split between clear and cloudy skies, Figure 4 illustrates the incongruity of the endeavour. It is rare to have a continuously clear sky day, and thus an invariance to atmospheric effects is required in training DNN models to mitigate the generation of undesired artifacts. In Pix2PixHD, this was observed through the generation of clear skies with artifacts at the localities of recurrent cloud formations. Figure 5 illustrates that SkyNet has invariance and successfully learned to generate clear skies. As illustrated, SkyNet produces some visual artifacts during sunset/sunrise, which are largely due to temporal class imbalances within HDRDB (fewer image samples of sunsets/sunrises). In future work, this will be corrected by training with stratified sampling of HDRDB, with the assumption that sunrises and sunsets can be considered equivalents. Though solar placement obviously differs, this assumption is presumed acceptable for negligible variations (e.g. in atmospheric composition and densities). This assumption is a subject of debate in photography, but it has been shown that visual differences are negligible and human subject are unable to reliably distinguish the subtle differences [22].

A.3 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 6c. This format is sub-optimal for DNN modelling, as the $360^\circ M \times N$ 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 angular format as defined by the solid angles in Figure 6b and shown in Figure 6d.

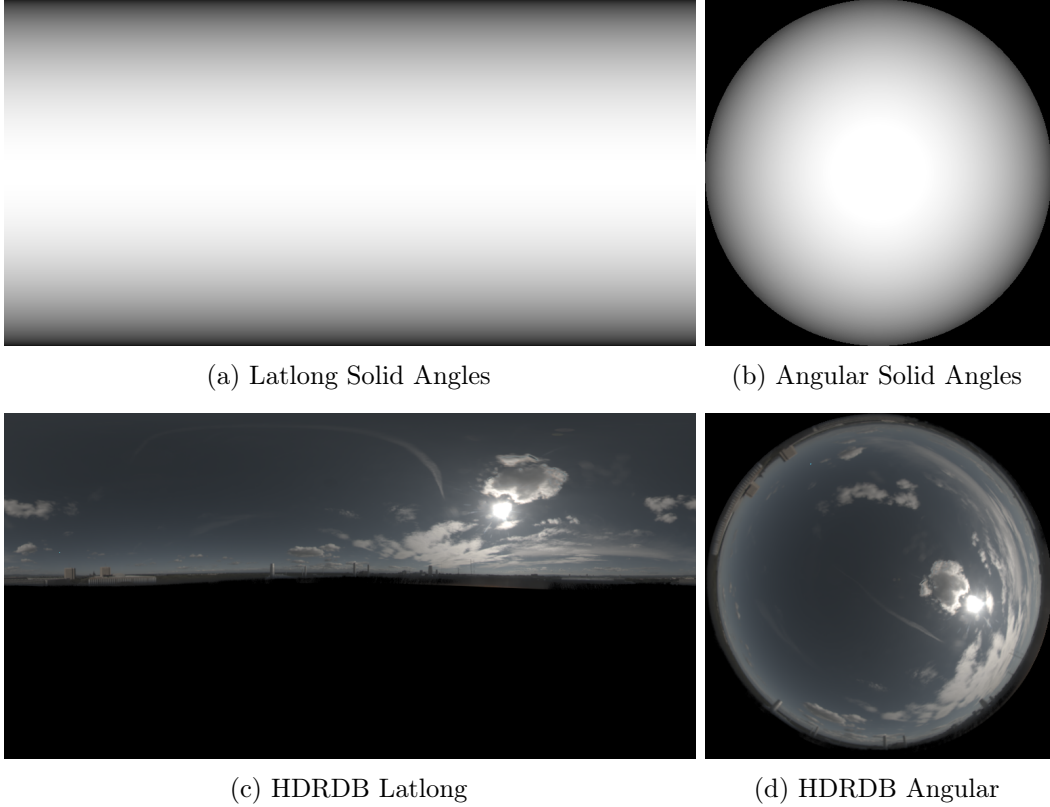


Figure 6: Latlong ($M \times N$) and angular ($N \times N$) environment maps

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 angle (φ) and azimuth angle (θ) as illustrated in Figure 7.

The image plane coordinate system denoted in terms of $+X$ columns and $+Y$ rows of pixels in Figure 7, illustrates the inequality the in skydome ‘surface area’ sampled by each pixel. As illustrated in Figures 6a and 6b, 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 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 $M \times N$ latlong skydomes by a factor of 2^x before transforming to angular representation. In experimentation this approach provided the best retention of radiant energy, with the latlong environment map shown in Figure 6c 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 angular representation as shown in Figure 6d.

A.4 Parametric Clear Skies

Figure 8 illustrates three parametric clear skies rendered using the Blender Creation Suite [23] with skies generated as scene backgrounds for fisheye lens ‘capture’. Where possible, all clear sky parameters were kept constant.




Figure 7: Definition of a skyangular skydome for solar angle (φ) and Azimuth angle (θ)

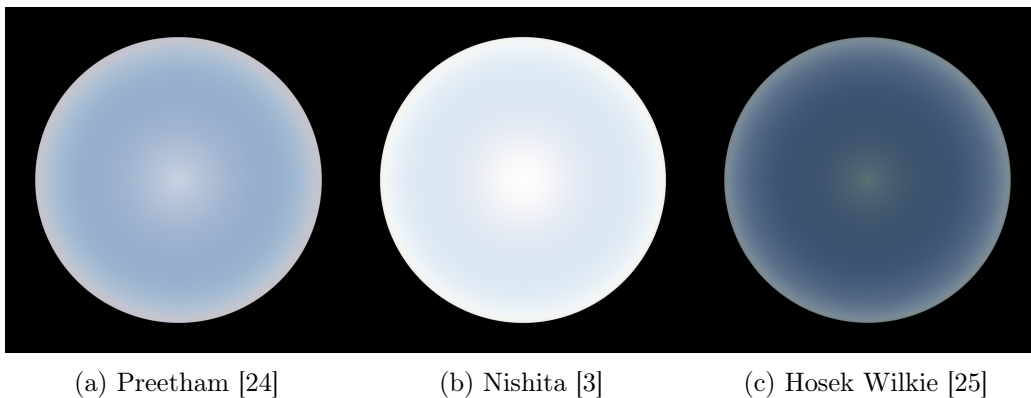


Figure 8: Blender Parametric Clear Skies

A.5 Validation of Dynamic Range

A.5.1 Tone-Mapping. Figure 9 illustrates the high dynamic range exhibited by clear skies (Figure 9a), and the impact of tone-mapping (Figure 9b) in the RGB color space and luminance space Y . We compress the dynamic range of each HDRI I_{in} using global logarithmic tone-mapping $I_{out} = \log_2(I_{in} + 1)$. This effectively reduces the dynamic range by an approximate factor of $\times 10,000$, which can be inverted via $I_{in} = 2^{(I_{out})} - 1$. Our model is 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 in 3D rendering engines.

We do not clip the dynamic range of the image I_{out} to the range $[0, 1]$ to mitigate the loss in the reconstructed dynamic range of Image I_{in} as shown in Figure 9b. For visualization, images illustrated in this paper are Gamma 2.2 tone-mapped and clipped to the range $[0, 1]$.

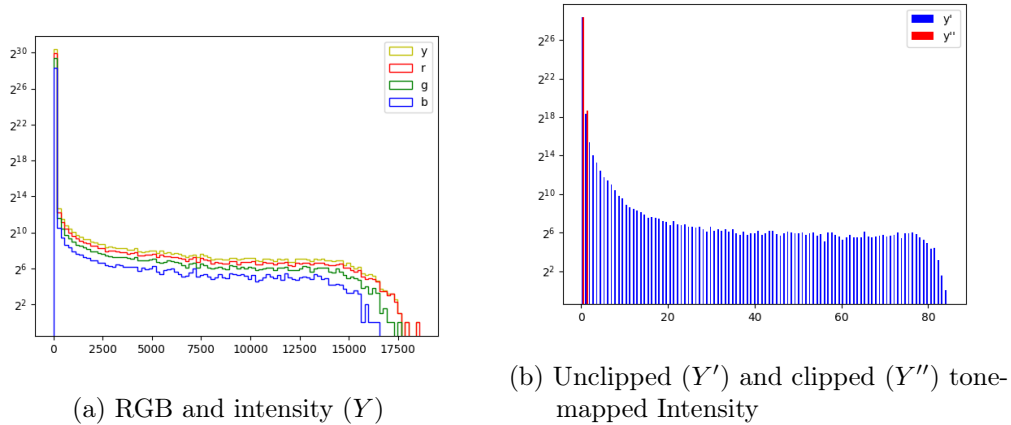


Figure 9: Histograms of HDR Skydomes from the Laval HDRDB illustrating the number of pixels (y-axis) against their intensity (x-axis) for a set of 325 clear sky samples from June 14th, 2015. Figure 9b illustrates the impact of tone-mapping, with and without clipping of the tone-mapped range

A.5.2 Histogram Metrics. To validate HDRI generated by SkyNet, metrics were implemented to compare the histograms with paired ground-truth images from HDRDB. To provide a comparison across a day-cycle, a novel set of 393 clear sky samples collected between 6:45AM and 7:53PM May 14th, 2015 were selected. For each image pair, the luminance spectrum (grayscale) was discretized into density distribution histograms for the range $[0, 100]$ with 1,000 bins and compared via six metrics. As illustrated in Figure 10, our method generates faithful representations of ground truth images throughout the day-cycle. This is demonstrated by stable high-correlation and low χ^2 , χ_{alt}^2 , EMD and KLD scores. Please see below for further definition of these metrics.

Intersection [26]:

$$d(H_1, H_2) = \sum_I \min(H_1(I), H_2(I)) \quad (1)$$

Correlation [26]:

$$d(H_1, H_2) = \frac{\sum_I (H_1(I) - \bar{H}_1)(H_2(I) - \bar{H}_2)}{\sqrt{\sum_I (H_1(I) - \bar{H}_1)^2 \sum_I (H_2(I) - \bar{H}_2)^2}} \quad (2)$$


where

$$\bar{H}_k = \frac{1}{N} \sum_J H_k(J) \quad (3)$$

and N is a total number of histogram bins.


Chi-Square (χ^2) [26]:

$$d(H_1, H_2) = \sum_I \frac{(H_1(I) - H_2(I))^2}{H_1(I)} \quad (4)$$




(a) Histogram Intersection \uparrow

(b) Histogram Correlation \uparrow



(c) $\chi^2 \downarrow$

(d) $\chi^2_{alt} \downarrow$



(e) Wasserstein Distance \downarrow
(Earth Mover's Distance; EMD)

(f) Kullback Leibler Dist (KLD) \downarrow

Figure 10: Histogram Comparison Metrics

Alternative Chi-Square (χ_{alt}^2) [27, 26]:

$$d(H_1, H_2) = 2 * \sum_I \frac{(H_1(I) - H_2(I))^2}{H_1(I) + H_2(I)} \quad (5)$$

Wasserstein Distance (Earth Mover's Distance; EMD) [28, 29]:

$$L_1(u, v) = \left(\inf_{\pi \in \Gamma(u, v)} \int_{\mathbb{R} \times \mathbb{R}} \|X - Y\| d\pi(x, y) \right) \quad (6)$$

where $\Gamma(u, v)$ is the set of all joint probability measures on $\mathbb{R} \times \mathbb{R}$ whose marginals are u, v on the first and second factors respectively. For Cumulative Distribution Functions (CDF) as used in this work, EMD can be re-written as:

$$L_1(u, v) = \sum \|U - V\| \quad (7)$$

Kullback-Leibler divergence (KLD) [26]:

$$d(H_1, H_2) = \sum_I H_1(I) \log \left(\frac{H_1(I)}{H_2(I)} \right)$$