

Illumination Space: A Feature Space for Radiance Maps

A. Chalmers¹, T. Zickler², and T. Rhee¹

¹Computational Media Innovation Centre, Victoria University of Wellington, New Zealand

²Harvard University, United States of America

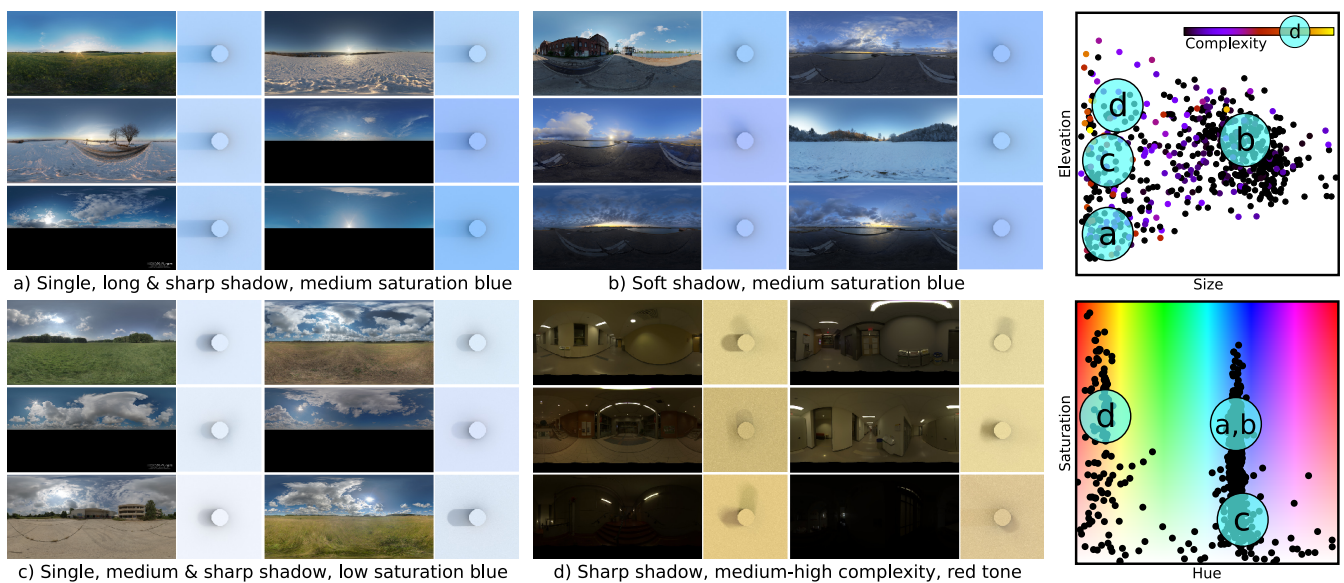


Figure 1: Radiance maps within the same cluster share similar illumination properties (a, b, c, and d). The features are visualized in the two scatterplots, which we call the ‘illumination space’, on the right.

Abstract

Radiance maps (RM) are used for capturing the lighting properties of real-world environments. Databases of RMs are useful for various rendering applications such as Look Development, live action composition, mixed reality, and machine learning. Such databases are not useful if they cannot be organized in a meaningful way. To address this, we introduce the illumination space, a feature space that arranges RM databases based on illumination properties. We avoid manual labeling by automatically extracting features from an RM that provides a concise and semantically meaningful representation of its typical lighting effects. This is made possible with the following contributions: a method to automatically extract a small set of dominant and ambient lighting properties from RMs, and a low-dimensional (5D) light feature vector summarizing these properties to form the illumination space. Our method is motivated by how the RM illuminates the scene as opposed to describing the textural content of the RM.

1. Introduction

Radiance maps (RMs) are high dynamic range (HDR) 360° images that are suitable for storing real-world lighting. Databases of RMs are now readily available in post-production studios and online ser-

vices, providing a large amount of high fidelity illumination data. Current machine learning methods use RM databases for unsupervised learning tasks [GSY*17, HGSH*17]. However, an unlabeled database of RMs is not useful in many other lighting applications that require searching or browsing the database (e.g., Look De-

velopment, live action composition, mixed reality rendering, and supervised learning). Furthermore, the labels should be concise to avoid sparsity introduced by high-dimensional data (i.e. the curse of dimensionality). An intuitive, low-dimensional representation that describes lighting, not texture, is required for many of these applications. Describing an RM in concise form is challenging since each RM contains a variety of complex lighting properties including shading, tone, shadows, and glossy highlights.

Current models either do not provide enough semantically meaningful parameters (e.g., Hosek-Wilkie [HW12] does not cover overcast skies or indoor lighting), or provide too many unintuitive parameters (e.g., spherical harmonics [RH01]). To overcome these limitations, we formulate our problem to utilize spherical Gaussians (SG) [TS06, XSD*13] and diffuse shading, which previously has not been used to categorize RM databases in a low-dimensional and semantically meaningful way. We aim to capture properties that are observable in the rendered scene rather than the textural content of the RM image itself. To obtain a low-dimensional representation, we introduce a dominant light model (DLM) which uses SGs to describe the most visually salient lighting effects as observed in the rendered scene. Diffuse shading is then used to capture semantically meaningful, low-dimensional tonal properties.

To develop a low-dimensional representation, we follow an empirical approach of making careful design decisions and observing their effects in large RM databases. Numerous iterations of this design process lead us to five intuitive lighting properties of RMs, split into two categories, *dominant* and *ambient* lighting: the (1) **size**, (2) **elevation angle**, and (3) **azimuth angle** of the RM's most dominant area light sources, as well as the diffuse (4) **hue** and (5) **saturation** of the RM. The dominant lighting property set (1, 2, 3) can occur multiple times for a given RM (one set for each bright light in the RM). To create a low-dimensional feature space, these properties are converted into a strictly 5D feature vector. This defines the '*illumination space*' - a semantically meaningful, low-dimensional feature space (see Figure 1) that can enable browsing and searching RM databases. Our contributions are summarized as follows:

- We define and extract a small set of dominant and ambient lighting properties from RMs using the DLM.
- From the extracted lighting properties, we introduce a 5D *light feature* vector that encodes RMs in a low-dimensional, semantically meaningful feature space - the '*illumination space*'.

2. Related Work

2.1. Textural Features for Sky Images

Manual labels applied to images (or parts of images) can guide browsing and searching in many applications [JGJJ*06, LHE*07, LRT*14]. Labels extracted from the metadata attached to images at capture time can also be used. However, these sorts of labels have limited expressibility and are not always available, so features computed from image content—such as GIST [OT01], geometric context [HEH05], and spatial pyramids [LSP06]—are often used instead. When they are applied to an RM, we refer to these as “textural features” because they are computed from the contrast patterns of the RM itself, instead of the lighting patterns

that the RM induces in a rendered scene. Tao et al. [TYS09] use textural features and supervised learning techniques to develop an interactive search system for finding sky photographs. Ono et al. [ODY11] extract four textural features to characterize the images in a database, with ongoing improvements for sky searching and generation [MF12, LRT*14]. Similarly, Chalmers et al. [CLH*14] use textural features to enable artists to find backdrops with various amounts and types of clouds. A substantial limitation of textural features is that they often do not adequately encode an RM's lighting effects. This is what we aim to address in this paper.

2.2. Perceptual Embeddings

There is a long history of creating low-dimensional and perceptually meaningful embeddings of materials for digital design and editing tasks. The 1976 CIELAB representation of color is an early example. More recent examples include perceptual embeddings of reflectance [PFG00, WAKB09, SGM*16] and translucency [GXZ*13, PRJ*13]. We follow this same general strategy, albeit for studying illumination—and RMs in particular—instead of materials. Previous encodings for RMs have considered dimensionality reduction without perceptual uniformity [KSH*14], or have relied on the manual attachment of tags [BBH*08], such as sun position and color, which is difficult to extend to large databases and the many possible attributes of appearance.

2.3. Sampling and Compression for RMs

Frequency space representations have been used to compress RMs down to as few as nine [RH01, SKS02] or hundreds [NRH03] coefficients with varying degrees of fidelity. Our 5D light feature vector instead describes both high and low frequency lighting in a semantically meaningful way with fewer parameters. Sampling algorithms sample areas of interest in the RM [Deb06, VD09, ARBJ03, LPG13], however, these sampling methods aim to oversample the area of interest rather than sample the light source once with a low-dimensional descriptor.

3. Dominant Light Definition

Before introducing the 5D light feature vector, we first define how we obtain dominant lighting properties, which we then use to encode into the dominant light features. A dominant light in an RM is a local region of the lighting sphere $\Omega \subset \mathbb{S}^2$ with the property that the region has a substantial impact on high-frequency lighting effects in a scene, including highlights and cast shadows. An RM may have multiple dominant lights, and we use Ω_j to denote the angular support of each one. We assume light Ω_j is given through a light detection scheme; we use statistical based thresholding defined by Rhee et al. [RPAC17]. The remainder of this section introduces the dominant light model (DLM) for summarizing the lighting effects of each Ω_j , and the procedure of fitting the DLM to each Ω_j . Section 4.1 shows how the fitted parameters of this model are used to compute the dominant light features of the illumination space.

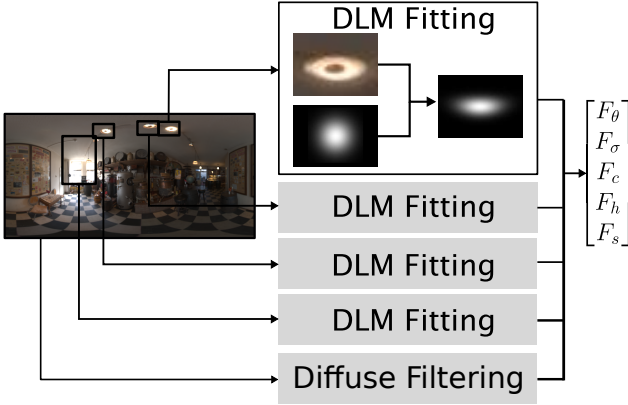


Figure 2: Feature extraction process. Given an input RM, each dominant light is fit with the DLM. The RM itself is also diffusely filtered. Then the elevation θ and size σ properties from each DLM are aggregated into three dominant lighting features F_θ , F_σ , and F_c . The hue and saturation features F_h and F_s are computed from the diffusely filtered RM.

3.1. Dominant Light Model (DLM)

To extract the elevation and size of a dominant light Ω_j , the first step is to fit a parametric model to each Ω_j . While we are only interested in a lower dimensional representation that Spherical Gaussians (SG) [TS06] provide, we use Anisotropic Spherical Gaussians (ASG) [XSD*13] since they often provide a better fit. We then reduce the ASG back into a SG. The ASG is defined as a function given a unit direction \mathbf{v} and light intensity a :

$$G(\mathbf{v}, a; [\mathbf{x}, \mathbf{y}, \mathbf{z}], [\sigma_x, \sigma_y]) = a \cdot \max(\mathbf{v} \cdot \mathbf{z}, 0) \cdot e^{-\sigma_x(\mathbf{v} \cdot \mathbf{x})^2 - \sigma_y(\mathbf{v} \cdot \mathbf{y})^2} \quad (1)$$

The vectors \mathbf{x} , \mathbf{y} , $\mathbf{z} \in \mathbb{S}^2$ are the lobe, tangent and bi-tangent axes respectively, forming an orthogonal basis of R^3 , with \mathbf{z} being the ‘direction’ of the dominant light relative to objects in a rendered scene. The parameters σ_x and σ_y define the size of the Gaussian for the x and y axes respectively (where $\sigma_x, \sigma_y > 0$), and a is the light intensity. The ASG is reduced into a SG by averaging σ_x and σ_y .

3.2. DLM Fitting

The process of fitting $G(\mathbf{v}, a; \mathbf{p})$ (where \mathbf{p} are the parameters of an DLM) to a dominant light source $f(\mathbf{v})$ provides an automatic parameterization of real-world lighting. From the fitted $G(\mathbf{v}, a; \mathbf{p})$, we extract the light features that describe the real-world illumination properties. To fit $G(\mathbf{v}, a; \mathbf{p})$ to $f(\mathbf{v})$, we first initialize, then optimize. The input parameters \mathbf{v} and a are given from the direction of the local maxima and the intensity of the pixel in the RM from that direction. The remaining parameters \mathbf{p} from equation (1) are given an initial approximation (as well as upper and lower bounds) and are optimized. σ_x and σ_y are initialized to be radially symmetric with a solid angle of 0.5334° (the solid angle of the sun as observed from Earth), and a lower and upper limit of 0.5334° and 90° respectively. The rotation of the DLM is initialized to 0 with a lower and upper limit of 0° and 360° respectively. Given the initialization parameters and boundary conditions, the parameters \mathbf{p} are

optimized using the dog-leg trust-region algorithm [VL04] with the objective function

$$f(l, g) = \frac{1}{N} \sum_i w_i (l_i - g_i) \quad (2)$$

where l is the RM and the g is the DLM approximation of each light in the RM. N are the number of pixels in the RM, i is each pixel, and w is the pixel’s corresponding solid angle weight to account for the distortion at the poles of an equirectangular image.

4. Light Feature Vector

We denote the elevation, azimuth, and size of a given dominant light as θ , ϕ , and σ respectively. From equation 1, θ and ϕ are the angular components of \mathbf{z} , while σ is the average of σ_x and σ_y . An RM may contain multiple dominant lights, thus an RM may contain multiple sets of these properties. To obtain a low-dimensional embedding, we describe a way of aggregating these dominant lighting properties into strictly three dominant light features: F_σ , F_θ , and F_c . The first two features describe the size and elevation, whereas the third takes into account the impact of multiple dominant lights, which describe as “complexity” or “angular spread”. In addition to these dominant light features, we define the features F_s and F_h that describe the hue and saturation of the RM’s diffuse shading. An overview of this process is shown in Figure 2.

4.1. Dominant Light Features

After each dominant light Ω_j in an RM is fit by an instance of the DLM, the features of the RM are computed from the fitted DLMs. In the case that an RM contains more than one dominant light, there will be multiple DLMs producing $N \times 2$ light properties (σ, θ) for each of the N dominant lights. We identify the dominant light with the highest maximum intensity value as the *primary light*. The remaining $N - 1$ dominant lights are *secondary lights*. The elevation and size features (F_θ, F_σ) are taken directly from the primary light’s DLM θ and σ values. The primary light is chosen as it gives the user an intuitive handle on the most visually striking effect the RM has on the rendered scene. We then aggregate the secondary light DLMs into a single feature F_c by considering the azimuthal relationship between the primary and secondary lights. To do this, we first define a weighted azimuthal distance as follows:

$$h(\phi_p, \phi_j) = \frac{|\phi_p - \phi_j|}{\pi} \cdot \log_{11} \left(\frac{a_j}{\sigma_j} \cdot \frac{1}{I} + 1 \right), \quad (3)$$

which takes the normalized azimuth distance between the primary light’s ϕ_p and the secondary light’s ϕ_j and applies a logarithmic scaling factor. The scaling factor is a way of measuring how visibly apparent the light is by taking the light’s mass $\frac{a_j}{\sigma_j}$ and dividing it by the total intensity of the RM I . Since this is a function of light intensity, we apply logarithmic scaling to be consistent with human perception. The overall effect of the scaling factor is to filter out secondary lights that are not as visually striking in their azimuthal



Figure 3: Querying the database with the input (ground truth), we compare our result with [Karsch14]. Our method maintains similarities with the ground truth in both light (RM, left) and illumination properties (rendered scene, right).




Figure 4: Examples of compositing into a photograph, where the RM was selected by a user using the illumination space. Inserted objects are: (left) Stanford Dragon and Bunny, (center) Stanford Asian Dragon, and (right) cylinder.

spread of light/shadow. Finally, the feature is computed by taking the maximum value across all secondary dominant lights:

$$F_c = \max(h(\phi_p, \phi_j) : j = 1, \dots, N). \quad (4)$$

A multiplicative sum across using all secondary lights could be considered in order to ensure all dominant lights contributed to the feature. However, its physical meaning becomes less obvious. By using the max function, we can observe that a pair of dominant lights in an RM that were azimuthal opposite, with equal intensity, would become the highest possible azimuthal spread F_c . By introducing the RM’s intensity I , the feature scales for each RM, giving those RM’s with relatively brighter secondary lights a higher complexity value.

4.2. Ambient Light Features

The ambient illumination properties from an RM plays an important role in determining the mood and tone of a rendered image. We describe this effect with two features: the diffuse hue and saturation (F_h and F_s respectively). There are a number of possibilities

for representing the overall color of an RM. It is intuitive if there is a 1-to-1 correspondence between the features and an exact matching color in the illuminated scene. For this reason, the features are set to the color of the diffusely filtered RM in the up direction. This corresponds to the color of a diffuse plane parallel with the floor.

5. Results

Using the light features, we are able to construct the *illumination space*, a 5D feature space for RMs which is compact and intuitive. Querying the database runs at interactive rates using optimized KNN algorithms [ML14]. Search-based tasks uses an input to query the RM database in order to find similar illumination. We can simply project a new RM into the database using the same feature extraction methods. As such, an example RM provided from artists (e.g., captured on set) can be used to find other similar illumination (Figure 3 shows examples of using an RM to query the database to find another similar RM). To demonstrate the utility of the illumination space for rendering tasks, Figure 4 shows examples of compositing synthetic objects into photographs. Figure 5, 6, and 7 show the examples of gradually changing the parameters along light size (shadow softness), elevation (shadow length), and

the azimuthal spread features (F_σ , F_θ , and F_c respectively). We observe that the RMs often appear in one of two color distributions (the peaks of hue in Figure 1). We visualize the cool to warm transition in one dimension in Figure 8.

6. Conclusion

This paper presents a set of light features that describe the illumination properties of an RM. The features themselves are concise and intuitive, presented as a 5D feature vector, describing both high and low frequency details. The features have a direct correspondence between the light source and its impact on the rendered scene. By introducing a dominant light model, fitting algorithm, and lights features, it is possible to automatically parameterize a large database of RMs into low dimensional features. From these features, we are able to construct a lighting-based feature space - the illumination space - that arranges RM databases. Future work can consider methods of overcoming sparsity in RM databases (e.g., interpolation), which can make it difficult to navigate the space if few samples exist.

References

- [ARB03] AGARWAL S., RAMAMOORTHY R., BELONGIE S., JENSEN H. W.: Structured importance sampling of environment maps. *ACM Transaction on Graphics* (2003). 2
- [BBH*08] BLOCH C., BAUER C., HEISTERBERGE V., REDDICK G., CHRIS HUF: Smart-ibl. <http://www.hdrlabs.com/sibl/>, 2008. Accessed: 30-06-2020. 2
- [CLH*14] CHALMERS A., LEWIS J., HILLMAN P., TAIT C., RHEE T.: Sky Browser: Search for HDR Sky Maps. In *Pacific Graphics* (2014), The Eurographics Association. 2
- [Deb06] DEBEVEC P.: A median cut algorithm for light probe sampling. In *ACM SIGGRAPH 2006 Courses* (2006), ACM. 2
- [GSY*17] GARDNER M., SUNKAVALLI K., YUMER E., SHEN X., GAMBARETTO E., GAGNÉ C., LALONDE J.: Learning to predict indoor illumination from a single image. *ACM Transactions on Graphics (SIGGRAPH Asia)* (2017). 1
- [GXZ*13] GKIIOULEKAS I., XIAO B., ZHAO S., ADELSON E. H., ZICKLER T., BALA K.: Understanding the role of phase function in translucent appearance. *ACM Trans. Graph.* 32, 5 (Oct. 2013), 147:1–147:19. 2
- [HEH05] HOIEM D., EFROS A. A., HEBERT M.: Geometric context from a single image. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on* (2005), vol. 1, IEEE, pp. 654–661. 2
- [HGS*17] HOLD-GEOFFROY Y., SUNKAVALLI K., HADAP S., GAMBARETTO E., LALONDE J.-F.: Deep outdoor illumination estimation. In *IEEE International Conference on Computer Vision and Pattern Recognition* (2017). 1
- [HW12] HOSEK L., WILKIE A.: An analytic model for full spectral skydome radiance. *ACM Trans. Graph.* (2012). 2
- [JGJ*06] JOHNSON M. B., G. J. S., J. A., O. KWATRA V., CIPOLLA R.: Semantic photo synthesis. *Computer Graphics Forum* 25, 2 (2006), 407–412. 2
- [KSH*14] KARSCH K., SUNKAVALLI K., HADAP S., CARR N., JIN H., FONTE R., SITTIG M., FORSYTH D.: Automatic scene inference for 3d object compositing. *ACM Trans. Graph.* (2014). 2
- [LHE*07] LALONDE J.-F., HOIEM D., EFROS A. A., ROTHER C., WINN J., CRIMINISI A.: Photo clip art. *ACM Trans. Graph.* 26, 3 (July 2007). 2
- [LPG13] LU H., PACANOWSKI R., GRANIER X.: Real-time importance sampling of dynamic environment maps. In *EG 2013 - Short Papers* (2013). 2
- [LRT*14] LAFFONT P.-Y., REN Z., TAO X., QIAN C., HAYS J.: Transient attributes for high-level understanding and editing of outdoor scenes. *ACM Transactions on Graphics (proceedings of SIGGRAPH 2014)* 33, 4 (2014). 2
- [LSP06] LAZEBNIK S., SCHMID C., PONCE J.: Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *Computer vision and pattern recognition, 2006 IEEE computer society conference on* (2006), vol. 2, IEEE, pp. 2169–2178. 2
- [MF12] MITANI T., FUJISHIRO I.: Cosmicai: Generating sky backgrounds through content-based search and flexible composition. In *ACM SIGGRAPH 2012 Posters* (2012), ACM, pp. 52:1–52:1. 2
- [ML14] MUJA M., LOWE D. G.: Scalable nearest neighbor algorithms for high dimensional data. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 36 (2014). 4
- [NRH03] NG R., RAMAMOORTHY R., HANRAHAN P.: All-frequency shadows using non-linear wavelet lighting approximation. In *Proc. of ACM SIGGRAPH 2003* (2003), ACM. 2
- [ODY11] ONO A., DOBASHI Y., YAMAMOTO T.: A system for editing sky images using an image database. In *SIGGRAPH Asia 2011 Sketches* (2011), SA '11, ACM, pp. 38:1–38:2. 2
- [OT01] OLIVA A., TORRALBA A.: Modeling the shape of the scene: A holistic representation of the spatial envelope. *International Journal of Computer Vision* (2001). 2
- [PFG00] PELLACINI F., FERWERDA J. A., GREENBERG D. P.: Toward a psychophysically-based light reflection model for image synthesis. In *Proc. of SIGGRAPH '00* (2000). 2
- [PRJ*13] PAPAS M., REGG C., JAROSZ W., BICKEL B., JACKSON P., MATUSIK W., MARSCHNER S., GROSS M.: Fabricating translucent materials using continuous pigment mixtures. *ACM Trans. Graph.* (2013). 2
- [RH01] RAMAMOORTHY R., HANRAHAN P.: An efficient representation for irradiance environment maps. *SIGGRAPH '01*, ACM. 2
- [RPAC17] RHEE T., PETIKAM L., ALLEN B., CHALMERS A.: Mr360: Mixed reality rendering for 360 panoramic videos. *IEEE Transactions on Visualization & Computer Graphics*, 4 (2017), 1379–1388. 2
- [SGM*16] SERRANO A., GUTIERREZ D., MYSZKOWSKI K., SEIDEL H.-P., MASIA B.: An intuitive control space for material appearance. *ACM Trans. Graph.* (2016). 2
- [SKS02] SLOAN P.-P., KAUTZ J., SNYDER J.: Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *SIGGRAPH 2002*, ACM. 2
- [TS06] TSAI Y.-T., SHIH Z.-C.: All-frequency precomputed radiance transfer using spherical radial basis functions and clustered tensor approximation. *Transactions on Graphics* (2006). 2, 3
- [TYS09] TAO L., YUAN L., SUN J.: Skyfinder: Attribute-based sky image search. In *Proc. of ACM SIGGRAPH 2009* (2009), ACM, pp. 68:1–68:5. 2
- [VD09] VIRIYOTHAI K., DEBEVEC P.: Variance minimization light probe sampling. In *SIGGRAPH'09: Posters* (2009), ACM, p. 92. 2
- [VL04] VOGLIS C., LAGARIS I.: A rectangular trust region dogleg approach for unconstrained and bound constrained nonlinear optimization. In *WSEAS Conference* (2004), pp. 17–19. 3
- [WAKB09] WILLS J., AGARWAL S., KRIEGMAN D., BELONGIE S.: Toward a perceptual space for gloss. *ACM Trans. Graph.* 28, 4 (Sept. 2009), 103:1–103:15. 2
- [XSD*13] XU K., SUN W.-L., DONG Z., ZHAO D.-Y., WU R.-D., HU S.-M.: Anisotropic spherical gaussians. *ACM Transactions on Graphics* 32, 6 (2013), 209:1–209:11. 2, 3

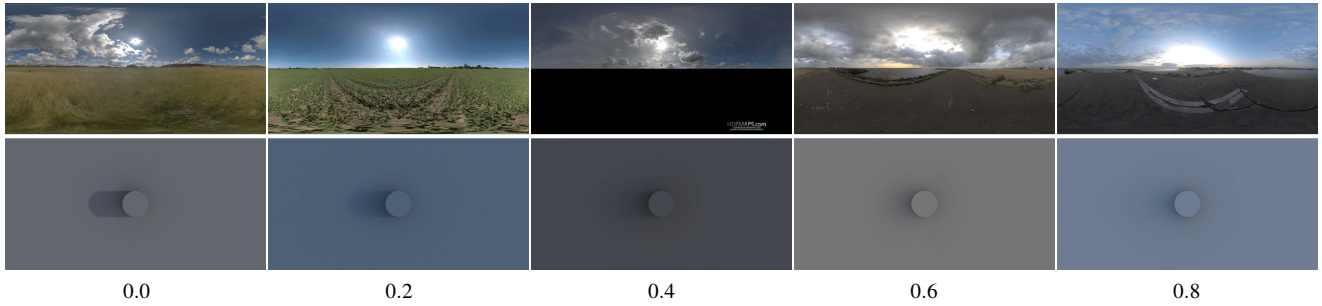


Figure 5: Feature F_{σ} , we show the transition from small to large area lighting (sharp to soft shadows).




Figure 6: Feature F_{θ} , we show the transition from low to high elevation lighting (long to short shadows).




Figure 7: Feature F_c , we show the transition from low to high azimuthal spread (complexity).




Figure 8: Feature F_h and F_s are the color features. It is common to find RMs with either red (indoor, sunset), blue (clear sunny sky) or desaturated colors (indoor, overcast). This corresponds to a cool/warm color distribution. We show transition from high saturation blue, desaturated, and high saturation red in one dimension.