

Deep Lighting Environment Map Estimation from Spherical Panoramas

Vasileios Gkitsas^{*1}, Nikolaos Zioulis^{*1,2}, Federico Alvarez², Dimitrios Zarpalas¹, and Petros Daras¹

¹Information Technologies Institute (ITI), Centre for Research and Technology Hellas (CERTH)

²Signals, Systems and Radiocommunications Department (SSRD), Universidad Politécnica de Madrid (UPM)

Abstract

Estimating a scene's lighting is a very important task when compositing synthetic content within real environments, with applications in mixed reality and post-production. In this work we present a data-driven model that estimates an HDR lighting environment map from a single LDR monocular spherical panorama. In addition to being a challenging and ill-posed problem, the lighting estimation task also suffers from a lack of facile illumination ground truth data, a fact that hinders the applicability of data-driven methods. We approach this problem differently, exploiting the availability of surface geometry to employ image-based relighting as a data generator and supervision mechanism. This relies on a global Lambertian assumption that helps us overcome issues related to pre-baked lighting. We relight our training data and complement the model's supervision with a photometric loss, enabled by a differentiable image-based relighting technique. Finally, since we predict spherical spectral coefficients, we show that by imposing a distribution prior on the predicted coefficients, we can greatly boost performance. Code and models available at vcl3d.github.io/DeepPanoramaLighting

1. Introduction

Compositing content from different domains (*i.e.* real and synthetic) has been an important part of post-production and visual effects. It has now also become very relevant due to the maturation of mixed reality technologies. These emerging technologies operate in between the real and virtual domains by embedding digitized or digital content in other media.

The realism of this composition depends on the positioning and lighting of the emplaced content, which is very demanding to accomplish, especially for monocular content as their estimation depends on solving ill-posed problems. Data-driven methods have advanced the state-of-the-art in

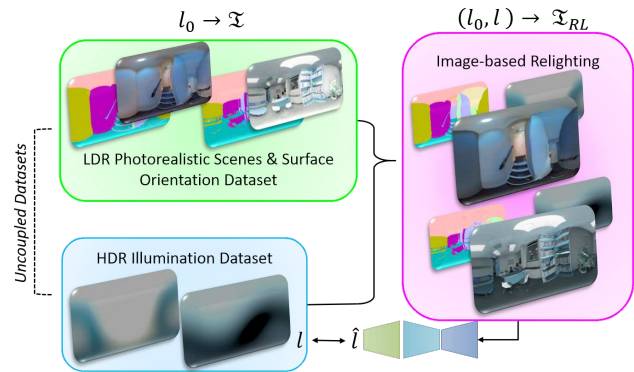


Figure 1: We use uncoupled datasets of HDR illumination maps [11] and photorealistic spherical panoramas [17] to train a spherical lighting environment map estimation model. While the former offer small content variance, they provide ground truth lighting l . While the latter offer multimodal data, they come with color content I that has pre-baked lighting l_0 . We combine these uncoupled datasets seamlessly using image-based relighting, producing relit images I_{RL} and condition our supervision solely on the imposed lighting l .

challenging tasks such as geometry estimation from single monocular images [8], and have even overcome the shortage of high quality data through self-supervision [37].

Nonetheless, this is far more challenging for illumination estimation, due to the very complex process associated with the coloring of each pixel in an image. Rarely – if not ever – do we see in the real world a scene to be exclusively lit from a single light source. In most cases illumination is composed of multiple sources including reflections, localized light sources (*e.g.* spot-lights) and/or broader area distributions (*e.g.* sunlight).

For the realistic relighting of embedded content, a high dynamic range (HDR) illumination needs to be estimated. This makes data collection problematic and impractical as predictions will typically be done on a low dynamic range (LDR) image. The Laval Indoor HDR Dataset [11] only contains a very limited – in terms of modern data-driven

^{*}Equal contribution

methods’ requirements – amount of coupled LDR and HDR spherical panoramas, capturing a scene’s global lighting conditions, with some being duplicate captures [36]. The Matterport3D dataset [4], used in [33] to generate illumination environment maps, contains saturated HDR images and misses ceilings, where most of the indoor lighting comes from. Further, the geometric warping used to register the LDR and HDR images resulted in high resolution artifacts. Another direction [12, 29, 22] would be to use synthetic scenes, but given the lack of realistic materials and textures, the model’s transferability to real world data would be limited. Other approaches rely on specialized capturing with expensive hardware [26, 35] or customized spherical objects [20].

Currently, the biggest challenge that data-driven illumination estimation methods need to overcome is the availability of data. Both [26] and [35] only captured 18 subjects, while [11] contains less than 2000 high quality LDR-HDR pairs. At the same time, Matterport3D offers a large amount of data but does not offer high quality HDR environment maps, with the derivative dataset of [33] coupling its generated HDR data with severely limited LDR samples.

Instead, we design a strategy that will exploit the strengths of all available datasets in order to train a model with the highest possible data variability – in terms of photometric content and illuminations. Our concept, as presented in Fig. 1, is based on image re-lighting (RL). We relight a single image \mathcal{I} with the lighting parameters ℓ through a rendering operation \mathcal{R} :

$$\mathcal{I}_{RL} = \mathcal{R}(\mathcal{I}, \mathcal{N}, \ell), \quad (1)$$

where \mathcal{N} represent dense per pixel metadata. Given estimated lighting parameters $\hat{\ell} = \mathbf{f}(\theta, \mathcal{I}_{RL})$ predicted by a model \mathbf{f} with parameters θ using \mathcal{I}_{RL} as input, we can train a model with uncoupled color and ground truth illumination information as long as we condition our supervision on ℓ alone. This is necessary as the image \mathcal{I} is a real world acquired image and has baked lighting into it. It is therefore the result of another physically based rendering operation that involved the original scene’s lighting parameters ℓ_o .

This concept crucially relies on a global Lambertian assumption [2]. This widely used assumption [5, 29] is a practical approach to alleviate the complexity of natural illumination. Under this assumption, we ensure that the original image \mathcal{I} – with pre-baked lighting – is a purely diffuse surface, effectively ignoring ℓ_o . In addition, under distant illumination, the relighting operation only depends on the surface’s global direction. As a result, our concept allows us to use high quality illumination maps (*i.e.* from [11]) and exploit the content plurality and multimodality of modern datasets (*i.e.* [4]) to train a global lighting map estimation model. Summarizing, in this work we train a model to estimate a lighting environment map from a single monocular

spherical panorama. Specifically, our contributions are:

- We use an efficient lighting representation and an image relighting rendering operation to synthesize randomized relit samples and thus, merge uncoupled datasets. We complement this with environment map blending to increasing the variability of our training data.
- As our image relighting operation is fully differentiable, we also leverage it on the supervision end to photometrically supervise our relit samples, enforcing a loss only on the imposed lighting.
- We further constrain our efficient spectral representation with a distribution prior that aids training offering a large performance gain.

2. Related Work

Traditionally, lighting estimation is a long-standing problem which is very challenging as it requires geometrical (*i.e.* 3D) data, camera (*i.e.* projection) characteristics, as well as material (*i.e.* attributes) metadata. Early progress relied on the acquisition of 3D shapes [25] and camera intrinsics information in order to estimate a scene’s lighting. Nonetheless, even modern methods [23] rely on accurate geometry information or on assumptions about it [24]. This line of work eventually matured into a single hypothesis estimation that explains the entirety of the scene (geometry and illumination) relying on natural scene statistics [1].

The seminal work of [6] first showed that it is possible to capture a scene’s HDR lighting by capturing differently exposed images of a mirrored metallic sphere that reflects light. Using such light probes has even enabled real-time video capture of HDR illumination at each position it is placed at within a scene. Multiple spherical light probes have also been used to estimate a scene’s global illumination [7]. More recently, a specialized shading probe was designed for capturing directly the scene’s shading for the purposes of mobile Augmented Reality (AR) [3]. Nevertheless, the physical placement of known 3D objects into the scenes is not always practical as the images whose lighting needs to be estimated might have been priory captured.

Even before the advent of deep models, the intractability of the problem of lighting estimation lend itself to data-driven priors [19] learned from millions of images. Yet, modern data-driven methods are much more suited to overcoming the challenges of ill-posed tasks. The main challenge that supervised data-driven methods need to overcome is the acquisition and/or estimation of supervisory data. Aiming to estimate outdoor illumination in [15], a synthetic sky model was fit to the captured panoramas in order to acquire the illumination parameters to regress. To address the shortage of data, [36] harvested a set of light sources from the Laval Indoor HDR dataset and used them to augment the dataset’s panoramas themselves, with the

goal of learning to predict the illumination from monocular images of known 3D objects. This was achieved by aligning the latent space of HDR panoramas using an autoencoder, and then predicting the illumination map from a perspective render of it, using lighting from the dataset’s panoramas in a supervised manner.

A custom data collection rig with three reflective spheres and a mobile phone was designed in [20] to acquire a very large dataset of paired images with reflective spheres imaged with the corresponding illumination. Then, an image-relighting supervision scheme was employed that relied on the photometric consistency between the sphere images collected by the capturing rig and the corresponding renders with the predicted illumination. Generalization to indoor and outdoor scenes equally, enabled mobile mixed reality.

Synthetic content can be used to bypass cumbersome and custom data collection like in [12], where the SunCG dataset [34] was used, which contains light positions. After generating scenes with randomized lighting parameters, a set of random light probes were positioned into these scenes and subsequently rendered into RGB-D cubemaps. A pre-trained network was used along with a discriminator to allow for better inter-domain generalization. It was trained to predict spatially varying light, conditioned by a specific location in each image. This was achieved by the addition of an extra modality (*i.e.* depth) offered by the synthetic data and a dual path (*i.e.* global and local) network architecture. Taking one step beyond, a parametric lighting representation was estimated in [10], with the goal being to more accurately model the localized nature of lighting. Per light information (*i.e.* position, direction, color and size) was extracted from the Laval Indoor HDR dataset which was also used for providing the color data as near field of view crops of the panoramas. Similar to [12], a pre-trained encoder was used to alleviate the lack of data and fine-tuned to the lighting estimation task. When used with geometry information, this parametric estimation allows for finer grained lighting results that respect visibility.

All the aforementioned methods estimate a scene’s global (*i.e.* omnidirectional) illumination from a single perspective image. The recent work of [33] addressed this through an image warping step, followed by an omnidirectional completion one, and then by a LDR to HDR estimation, finally encoding the HDR representation into an irradiance environment map. It was demonstrated that an end-to-end pipeline for these three tasks offered better illumination estimates. The aforementioned work also contributed a novel dataset by sampling from perspective HDR images into omnidirectional HDR maps at specific locales. In this way, their location conditioned predictions estimated localised spherical illumination from perspective images. However, the Matterport3D HDR captures are saturated and do not always include ceilings, meaning that important

lighting information may be missed. In addition, the geometric warping produces LDR and HDR images with severe high frequency artifacts which are not realistic. Nonetheless, training on a higher variety of data and relying on omnidirectional supervision for the completion task helped in producing high quality omnidirectional lighting estimates, albeit object localised.

In our work, we estimate the illumination directly from a spherical panorama which captures global scene information more accurately. We also use the original indoor scene images from Matterport3D instead of the distorted object localised ones. This means that we need to overcome the lack of paired lighting information.

3. Lighting Estimation via Relighting

Our concept is based on the works for precomputed radiance transfer [32] and irradiance maps [28]. These works use spherical harmonics (SH) representations to precompute intermediate function evaluations and use them during rendering. SH are a frequency space representation¹ of a function over the sphere, analogous to the Fourier transform. More details about SH in general and SH lighting can be found in [14], however, we will followingly present the necessary preliminaries.

Spherical Harmonics Diffuse Lighting: Given that SH form an orthonormal basis over the sphere, each spherical function can be represented as a linear combination of a set of basis function:

$$\hat{\mathcal{L}}(\rho) = \sum_{\rho}^{\Omega} L_m^l Y_m^l(\rho), \quad (2)$$

where $\rho = (\phi, \theta)$ are the spherical angular coordinates defined in the spherical image domain Ω , L_m^l is the coefficient of the Y_m^l SH basis function of degree (*i.e.* band) l and order m , spanning $-l \leq m \leq l$. The function $\hat{\mathcal{L}}$ is an approximation that depends on the depth of the order l , with smaller orders corresponding to lower frequencies. In our context, the approximated function $\hat{\mathcal{L}}$ represents the distant illumination. The SH coefficients L_m^l are calculated by

$$L_m^l = \frac{4\pi}{N} \sum_{\rho}^{\Omega} \mathcal{L}(\rho) Y_m^l(\rho) \quad (3)$$

with $N = w \times h$ being the total discrete elements sampled, corresponding to the area of the spherical image domain of width w and height h .

As shown in [28], only the first 3 orders (*i.e.* $l = [0, 1, 2]$, meaning 9 coefficients) are needed to accurately represent the surface’s irradiance (up to 99.2% [2]), given that irradiance smoothly varies with orientation. Consequently, SH

¹While generally defined on imaginary numbers, we will only be referring to real valued functions, and thus, our definition of SH corresponds to Real Spherical Harmonic functions.

coefficients are a very effective technique to compress direct illumination from distant sources. For the remainder of this document we will consider SH coefficients only up to the third order $0 \leq \ell \leq 2$. As presented in [28], the resulting irradiance \mathcal{E} map can be expressed as a function of the surface’s orientation, represented by the normal map \mathcal{N} , and the SH coefficients $\ell = \{L_m^l\}$, after lifting them into a symmetric 4×4 matrix representation:

$$\mathcal{E}(\mathcal{N}, \ell) = \eta(\mathcal{N})^T \begin{bmatrix} c_1 L_2^2 & c_1 L_{-2}^2 & c_1 L_1^2 & c_2 L_1^1 \\ c_1 L_{-2}^2 & -c_1 L_2^2 & c_1 L_{-1}^2 & c_2 L_{-1}^1 \\ c_1 L_1^2 & c_1 L_{-1}^2 & c_3 L_0^2 & c_2 L_0^1 \\ c_2 L_1^1 & c_2 L_{-1}^1 & c_2 L_0^1 & c_4 L_0^0 - c_5 L_0^2 \end{bmatrix} \eta(\mathcal{N}), \quad (4)$$

with $\eta(\mathcal{N}) = (\mathcal{N}_x, \mathcal{N}_y, \mathcal{N}_z, 1)$ being a homogeneous coordinates transformation operation. The symmetric matrix and constants $c_i, i \in [1, 4]$ are the result of expanding the rendering equation and associated spherical harmonic normalization. Overall, we find that SH lighting – Eq. (4) – relies on a matrix-vector and a vector-vector product, which is fully differentiable, as well as parallelizable and thus, very well suited for modern tensor-based data-driven methods. Also, the same applies to the SH projection – Eq. (3) – and reconstruction – Eq. (2) – functions.

Training with Uncoupled Datasets: While works like [11] train on a limited amount of HDR panoramas, our goal is to learn illumination information from LDR images, similar to more recent works [20, 33, 10, 12]. Nonetheless, we need to regress HDR illumination in order to be able to convincingly add synthetic objects into natural scenes. More recent datasets either do not offer HDR environment maps [20] or their data come with deficiencies [33]. Seeking to avoid the problems associated to synthetic data [12], instead, our approach is to blend different lighting conditions to randomly synthesize illumination maps. Considering two HDR images representing dense lighting conditions \mathcal{L}_a and \mathcal{L}_b we can blend them using a ratio $\lambda_{blend} \in [0, 1]$ to produce a blended lighting map $\mathcal{L}_{ab} = \lambda_{blend} \mathcal{L}_a + (1 - \lambda_{blend}) \mathcal{L}_b$. Then, using Eq. (3) we can project \mathcal{L}_{ab} to a set of SH coefficients $\ell_{ab} = \{L_m^l\}$ up to a predefined order l . A set of examples of blending 2 lighting maps and their combined relighting effect in contrast to the original relights are presented in Fig. 2.

Considering a spherical image \mathcal{I} and its corresponding, aligned normal map \mathcal{N} we can relight the image \mathcal{I} to produce a new relit one \mathcal{I}_{RL} with a set of low dimensional SH lighting parameters ℓ using Eq. (4). Accordingly, the rendering operation \mathcal{R} defined in Eq. (1) for each pixel becomes:

$$\mathcal{R}(\mathcal{I}, \mathcal{N}, \ell) := \mathcal{E}(\mathcal{N}, \ell) \mathcal{I}. \quad (5)$$

Therefore, we can render new images $\mathcal{I}_{RL} = \mathcal{R}(\mathcal{I}, \mathcal{N}, \ell)$ coupled with the newly imposed lighting conditions ℓ that will serve as our training data. Even though the original image contains pre-baked lighting, as presented in Fig. 3, the




Figure 2: The combined effect of blending 2 lighting maps with an even ratio. On the left is the original color image and associated normal map. Then, on the right, 2 different lighting probes follow with their rendered lighting on the original image on top. Finally on the rightmost, we show the blended lighting probe as well as the rendered original image with the combined lighting on top. Lighting blending increases the variety of the dataset and can also help in overcoming non realistic relights.

new lighting can greatly influence the resulting image and change its overall appearance (through the ambient term), as well as relative scene luminance and shading. As long as we properly regularize our supervision on this newly imposed lighting alone, a lighting estimation model can be trained under this scheme.

This allows us to exploit larger photorealistic datasets with higher content variance, that additionally offer surface orientation information, and simultaneously exploit the higher quality, but limited in terms of scale, HDR illumination datasets. The scale limitation is bypassed through the aforementioned blending process that enables the on-the-fly generation of a larger HDR illumination dataset with higher variability compared to approaches such as [10] that only use the dataset’s original lighting. More specifically, we blend HDR illumination maps from the Laval Indoor HDR dataset [11] and relight the realistic panoramas [17] that contains raytraced normal maps as well.

Deep Spherical Lighting Estimation Model: Our CNN model is trained to regress a compressed representation of a lighting environment map that can be used to evaluate the irradiance of elements rendered within a scene. The input to our model is a LDR panorama in equirectangular format representing a scene for which we want to estimate the lighting. Two sub-networks are combined in an end-to-end manner to achieve this as depicted in Fig. 4. Since we need to regress to HDR illumination, a LDR-to-HDR autoencoder (AE) CNN is first used to translate the LDR image to the HDR domain. We then use an encoder (E) to regress to the SH coefficients $\hat{\ell} \in \mathbb{R}^{9 \times 3}$ that represent a low frequency distant illumination (9 SH coefficients) for each of the 3 color channels (*i.e.* red, green, blue).

Each training sample comprises a color image $\mathcal{I}^{LDR} \in \mathbb{R}^{W \times H \times 3}$, a normal map $\mathcal{N} \in \mathbb{R}^{W \times H \times 3}$, and the lighting coefficients ℓ calculated after randomly sampling two HDR illumination maps and blending them, to eventually produce

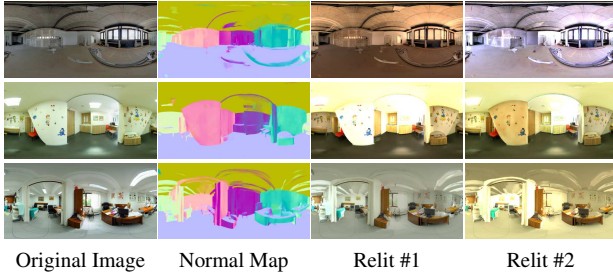


Figure 3: Exemplary relights of sample images, presented with their corresponding normal maps. The relit images’ lighting is highly modified and draws away from the original lighting. This allows us to largely overcome the pre-baked lighting while training, by considering the original image as a diffuse surface. Thus, we complement our supervision with a photometric supervision relighting loss that is only conditioned on the newly imposed light.

the network’s input which is a relit image \mathcal{I}_{RL}^{LDR} , using Eq. (5). In this way, we re-use and couple distinct, uncoupled data of color images supported by surface information, with other data of illumination maps. The relit LDR images gets translated by the LDR-to-HDR sub-network to an HDR image $\mathcal{I}_{RL}^{HDR} \in \mathbb{R}^{W \times H \times 3}$ which then gets encoded to the predicted SH coefficients $\hat{\ell}$.

The SH representation and the availability of surface information allows our supervision to be multi-faceted. First, we employ an L2 regression objective to the predicted coefficients, the coefficients loss:

$$\mathcal{L}_{SH} = \frac{1}{9 \times 3} \sum_{c=1}^3 \sum_{l=0}^2 \sum_{m=-2}^2 (L_{m,c}^l - \hat{L}_{m,c}^l)^2 \quad (6)$$

In addition, through Eq. (2) we can reconstruct the lighting maps \mathcal{L} and $\hat{\mathcal{L}}$ represented by the ground truth and estimated SH coefficients respectively. As a result, we further regularize our supervision with a denser L2 objective, the reconstruction loss:

$$\mathcal{L}_{RC} = \frac{1}{w \times h} \sum_{\rho} \mathcal{A}(\rho) \|\mathcal{L}(\rho) - \hat{\mathcal{L}}(\rho)\|_2^2, \quad (7)$$

where $\mathcal{A} \in \mathbb{R}^{W \times H \times 3}$ is the spherical attention mask used in [38]. In this way, pixels towards the equator are weighted more appropriately with respect to those near the poles that are sampled multiple times due to the equirectangular distortion and whose spherical area is smaller.

Finally, we can also compute a dense photometric consistency loss [13] between an image relit with the original lighting and that when relit using the model’s regressed output. As depicted in Fig 4, we convert the original (non-relit) LDR image to HDR, producing $\hat{\mathcal{I}}^{HDR}$. Using Eq. (5) and the normal map \mathcal{N} , we render this HDR image with the ground truth SH coefficients ℓ , producing $\hat{\mathcal{I}}_{GT}^{HDR}$, and with the estimated SH coefficients $\hat{\ell}$, creating $\hat{\mathcal{I}}_{RL}^{HDR}$. Our relighting loss is then formulated as:

$$\mathcal{L}_{RL} = \frac{1}{N} \sum_{\rho} \mathcal{A}(\alpha |\hat{\mathcal{G}}_{GT} - \hat{\mathcal{G}}_{RL}| + (1 - \alpha) SD(\hat{\mathcal{G}}_{GT}, \hat{\mathcal{G}}_{RL})), \quad (8)$$

where SD is the structural dissimilarity function, α a blending factor between SD and the L1 loss, while \mathcal{G} is the log transformed HDR image (superscripts and ρ indexing omitted for clarity). The L1 term in the log domain allows us to penalize big relighting errors, while the structural dissimilarity term penalizes shading discrepancies within regions of the image.

Therefore, our final weighted objective is:

$$\mathcal{L} = \lambda_{SH} \mathcal{L}_{SH} + \lambda_{RC} \mathcal{L}_{RC} + \lambda_{RL} \mathcal{L}_{RL} \quad (9)$$

It should be noted that relighting the LDR image requires a scaling of the SH coefficients to produce realistic and properly saturated relights. However, this is only needed for the input LDR alone, in order to modify its lighting. Supervision, regression and photometric alike, is facilitated by the unscaled – HDR – SH coefficients, allowing the network to learn HDR lighting from a single LDR image.

All prior works that regress SH coefficients directly operate on the model regressed values. In this way, the predicted values are in an unstructured form driven purely by the supervision. However, SH coefficients are spectral domain values and follow a structured distribution. The DC term (*i.e.* ambient or L_0^0) is the strongest coefficient, with the magnitude of the other frequencies dropping according to their order (*i.e.* $L_{\{-1,0,1\}}^1 \gg L_{\{-2,-1,0,1,2\}}^2$). We use this spectral distribution prior to enforce a structured SH prediction on each channel’s predicted coefficient vector $\hat{\ell}$ by applying a function $\tau(\hat{\ell})$, where:

$$\tau(\ell) = \|\ell\| \sigma(\ell). \quad (10)$$

The normalized exponential function σ (*i.e.* softmax) pushes lower values lower, and higher values higher, therefore easing the structuring of our predictions and aiding the supervision. Since it also normalizes to unity, we first extract the magnitude of the predicted coefficient vector, and reapply it after enforcing the distribution prior.

Finally, both the scaled coefficients used to relight the LDR input image, as well as the predicted coefficients are deringed. Ringing, also called Gibbs Phenomenon, was analyzed in [30] and relates to negative values during rasterization with the SH coefficients. Thus, we reduce ringing artifacts by multiplying the SH coefficients with a low pass filter [31], enforcing the coefficients per band to decrease smoothly to zero at some cut-off frequency. Our model is end-to-end trainable as all operations are differentiable.

4. Results

Implementation Details: Our experiments were run on a single machine with an i7 processor and a NVidia Titan X

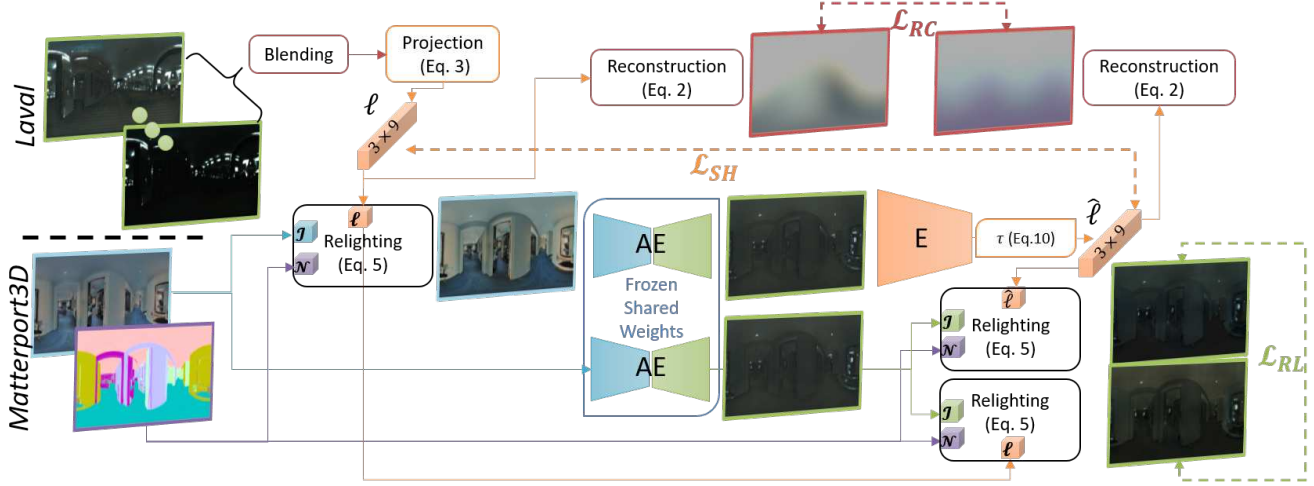


Figure 4: Our end-to-end relighting based supervision. An input LDR spherical panorama image \mathcal{I}^{LDR} is along with its corresponding normal map \mathcal{N} and randomly blended lighting parameters ℓ , in the form of SH coefficients, is relit using the rendering operation \mathcal{R} of Eq. (5). The relit LDR image \mathcal{I}_{RL}^{LDR} is translated to the HDR image \mathcal{I}_{RL}^{HDR} through the pre-trained (frozen weights) LDR-to-HDR autoencoder (AE). The result gets encoded by a lighting encoder E to the estimated lighting $\hat{\ell}$ which is passed through the spectral prior function τ (Eq. 10). The original LDR image also gets translated to the HDR domain, producing $\hat{\mathcal{I}}^{HDR}$. This gets rendered using using the aligned normal map \mathcal{N} two times, once with the original light ℓ and once with the predicted light $\hat{\ell}$. This way we synthesize $\hat{\mathcal{I}}_{GT}^{HDR}$ and $\hat{\mathcal{I}}_{RL}^{HDR}$ respectively which are used to photometrically supervise our lighting estimator (\mathcal{L}_{RL}). In addition, the ground truth and predicted SH coefficients are used to reconstruct lighting environments maps \mathcal{L} and $\hat{\mathcal{L}}$ respectively through Eq. (2). Additional losses are defined for the sparse SH coefficients (\mathcal{L}_{SH}) and the dense lighting map reconstructions (\mathcal{L}_{RC}). LDR images are with cyan border, HDR images and lighting parameters are with lime borders, while their low frequency reconstructions are with red, and normal maps are with violet borders.

GPU, using PyTorch [27]. All the models were trained for 10 epochs using the Adam optimizer [18] with its default parameters and an initial learning rate of 1×10^{-4} . The resolution $W \times H$ of the color images and their corresponding normal map is 512×256 . We generate randomized lighting coefficients from the Laval HDR Indoor Dataset [11], after blending the dataset’s HDR images in their original resolution and projecting them to SH using Eq. (3). However, when reconstructing them for the evaluation of \mathcal{L}_{RC} using Eq. (2) we reconstruct them in the training images’ resolution. Out of the 2235 images of the Laval dataset, we used 1682 for training and out of the remaining images, we kept 443 for evaluation purposes, while skipping 110 images due to very low brightness that led to very dark relit images. When relighting the LDR images, we scale the SH coefficients by 100 to align the dynamic ranges appropriately for LDR inputs.

The LDR-to-HDR autoencoder’s architecture was inspired by [9]. We use the dataset of [33] to train this sub-network, after resizing its coupled LDR and HDR images to 512×256 , and directly supervise it with a $L2$ reconstruction loss on the predicted HDR, and a $L2$ reconstruction loss on the reconstructed low frequency lighting environment map, which is the result of projecting the predicted HDR image to the SH basis and then reconstructing it from these SH coefficients. We use a learning rate of 1×10^{-3} and the Adam optimizer with its default parameters. We train this

network for 30 epochs without any weighting between the 2 losses. After convergence, we freeze the weights of the LDR-to-HDR model and use it to translate the input images to the HDR domain when training our lighting estimation encoder.

For the HDR lighting encoder, we base our model on the corresponding architecture of [15], using 7 convolutional and 2 fully connected layers with ELU activation functions except for the head of the encoder that regresses the output coefficients. When blending pairs of HDR environment maps from the Laval dataset, we randomly (uniform) sample two lighting maps and blend them using $\lambda_{blend} = 0.5$. This allows us to greatly increase the diversity of the estimated lightings. Our loss weights are $\lambda_{SH} = 0.01$, $\lambda_{RC} = 0.3$ and $\lambda_{RL} = 0.7$.

Quantitative Comparisons: Up to now, the state-of-the-art has focused on researching illumination estimation methods for traditional perspective images. On the other hand, our method estimates global lighting from omnidirectional images. Consequently, direct comparison is feasible with methods that do not rely on similar assumptions or learned perspective priors. We offer comparison results to SIRFS [1] which represents the state-of-the-art in learning free illumination estimation. We report the median scaled RMSE (m-RMSE) for the predicted lighting environment maps and the ground truth ones. Median scaling uses the median of the evaluated and ground truth signals to scale

the former to the value range of the latter. In this way we offer more meaningful comparisons between different methods by removing any scale ambiguity between their predictions which is required for numerous reasons. Methods like SIRFS that analytically solve the problem using LDR inputs, cannot easily regress to the correct HDR scale. Also, different datasets might offer HDR images at different scales, a prominent example being the saturated HDRs of Matterport3D. Finally, tone mapping, gamma correction, exposure times and color calibration are all different issues that influence the final compositing result. In the end, numeric accuracy might not always translate to high quality compositing, but instead, relatively correct estimates that better capture the light’s orientation and drop-off will surely produce more visually pleasing results. Finally, the reported m-RMSE uses spherical distortion weighting when evaluating the MSE, similar to [38], in order to reduce the metric skewing effect that equirectangular distortion introduces.

Table 1 presents the median scaled RMSE for our model and SIRFS on the test set of the Laval Indoor HDR dataset. The ground truth is generated by projecting the HDR images into the spherical harmonics basis using Eq. (3) and then reconstruct the ground truth lighting environment maps using Eq. (2). As both SIRFS and our models estimate SH coefficients, we reconstruct the corresponding dense environment maps and calculate m-RSME on these. Evidently, our model outperforms SIRFS, but more importantly it should be noted that the Laval images are *unseen* for our model, and that the estimated lighting maps never participated entirely during training due to blending.

Table 1 also presents results for a set of ablation experiments. Overall, we observe a close to 50% performance boost when imposing the spectral prior function τ on our predictions as indicated by the middle entries (without the prior), compared to the bottom entries (with the prior). Interestingly, relying purely on the relighting loss (the *only photo* variant with $\lambda_{SH} = 0, \lambda_{RC} = 0$) allows for proper model training. Further, training without the relighting loss (the *no photo* variant with $\lambda_{RL} = 0$) hurts performance, indicating its significance. At the same time, relying purely on the dense reconstructed signal loss (the *only dense* variant with $\lambda_{SH} = 0, \lambda_{RL} = 0$) showcases inferior performance even to SIRFS. More importantly, the discrepancy is a lot larger when not using the spectral prior function τ , whilst with it, its implicit regularisation does not allow for significant performance reduction.

Qualitative Results: Finally, we present qualitative results in Fig. 5 of synthetic object rendering into the spherical panoramas using our estimated lighting environment maps. We use the Mitsuba [16] ray-tracer engine and the Bunny and Armadillo models from the Stanford 3D dataset [21]. Each model is rendered with 512 samples and is presented with three different materials. We observe that the 0th order

Table 1: Median scaled RMSE results for the Laval test set.

Method	Prior	m-RMSE
SIRFS	No	0.0391
Ours (<i>full</i>)	No	0.0229
Ours (<i>only photo</i>)	No	0.0297
Ours (<i>no photo</i>)	No	0.0334
Ours (<i>only dense</i>)	No	0.0423
Ours (<i>full</i>)	Yes	0.0101
Ours (<i>only photo</i>)	Yes	0.0116
Ours (<i>no photo</i>)	Yes	0.0156
Ours (<i>only dense</i>)	Yes	0.0162

SH coefficient L_0^0 – the ambient lighting – is well estimated, as well as the 1st order $\{L_{-1}^1, L_0^1, L_1^1\}$ coefficients that encode directional lighting. Overall, our approach attains reasonable results even on completely unseen data harvested from the internet (first row).

5. Discussion

Lighting estimation is a very challenging and ill-posed problem. Instead of relying on direct supervision, recent data-driven works opt for photometric supervision by adding spherical objects during data collection [20], relying on synthetic data [12] or on lower quality warped data [34]. The most widely used HDR lighting dataset [11] used in some works [11, 10] only contains limited samples. Nevertheless, we showed that it is possible to learn lighting estimation in a supervised way with uncoupled data and thus, exploit the availability of high quality HDR lighting datasets in combination with larger and more diverse datasets that also contain surface information. Further, we showed that photometric supervision through image-based relighting is sufficient to drive a lighting estimator model. More importantly, we demonstrated that imposing the spectral distribution prior on the regressed SH coefficients greatly boosts model performance.

Nonetheless, our work only regresses up to the third order SH coefficients offering low frequency lighting environment map estimates. While a low frequency estimation behaves good metrically, it does not capture details, which is an important goal towards realism and is left as future work. Another issue is that without any geometrical knowledge it is not possible to embed visibility information which also detracts realistic relighting. Most other works also predict location-conditioned outputs while we offer global estimates, offering another direction for future work. Finally, our work is one of the first works to regress lighting from spherical panoramas, an emerging media type.

Acknowledgements: We acknowledge HW support by NVidia and financial support by the H2020 EC project Hyper360 (GA 761934).



Figure 5: Qualitative results for virtual object rendering in real scenes with the lighting estimated by our model. First seven rows are in-the-wild samples from **HDRI Haven** while last two rows are from the Laval test set. First six examples show perspective renders (viewport denoted within the panorama), while the last three rows show renders within the panoramas themselves. Three materials are used from left to right: a conductor (reflecting mirror), rough plastic (interior and exterior index of refraction of 1.9 and 1, respectively) and another conductor (gold).

References

- [1] Jonathan T Barron and Jitendra Malik. Shape, illumination, and reflectance from shading. *IEEE transactions on pattern analysis and machine intelligence*, 37(8):1670–1687, 2014. 2, 6
- [2] Ronen Basri and David W Jacobs. Lambertian reflectance and linear subspaces. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (2):218–233, 2003. 2, 3
- [3] Dan A Calian, Kenny Mitchell, Derek Nowrouzezahrai, and Jan Kautz. The shading probe: Fast appearance acquisition for mobile ar. In *SIGGRAPH Asia 2013 Technical Briefs*, page 20. ACM, 2013. 2
- [4] Angel Chang, Angela Dai, Thomas Allen Funkhouser, Maciej Halber, Matthias Niebner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. In *7th IEEE International Conference on 3D Vision, 3DV 2017*, pages 667–676. Institute of Electrical and Electronics Engineers Inc., 2018. 2
- [5] Dachuan Cheng, Jian Shi, Yanyun Chen, Xiaoming Deng, and Xiaopeng Zhang. Learning scene illumination by pairwise photos from rear and front mobile cameras. In *Computer Graphics Forum*, volume 37, pages 213–221. Wiley Online Library, 2018. 2
- [6] Paul Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '98*, pages 189–198, New York, NY, USA, 1998. ACM. 2
- [7] Paul Debevec, Chris Tchou, Andrew Gardner, Tim Hawkins, Charis Poullis, Jessi Stumpfel, Andrew Jones, Nathaniel Yun, Per Einarsson, Therese Lundgren, et al. Estimating surface reflectance properties of a complex scene under captured natural illumination. *Conditionally Accepted to ACM Transactions on Graphics*, 19, 2004. 2
- [8] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in neural information processing systems*, pages 2366–2374, 2014. 1
- [9] Gabriel Eilertsen, Joel Kronander, Gyorgy Denes, Rafał K. Mantiuk, and Jonas Unger. Hdr image reconstruction from a single exposure using deep cnns. *ArXiv*, abs/1710.07480, 2017. 6
- [10] Marc-André Gardner, Yannick Hold-Geoffroy, Kalyan Sunkavalli, Christian Gagné, and Jean-François Lalonde. Deep parametric indoor lighting estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7175–7183, 2019. 3, 4, 7
- [11] Marc-André Gardner, Kalyan Sunkavalli, Ersin Yumer, Xiaohui Shen, Emiliano Gambaretto, Christian Gagné, and Jean-François Lalonde. Learning to predict indoor illumination from a single image. *ACM Transactions on Graphics (TOG)*, 36(6):176, 2017. 1, 2, 4, 6, 7
- [12] Mathieu Garon, Kalyan Sunkavalli, Sunil Hadap, Nathan Carr, and Jean-François Lalonde. Fast spatially-varying indoor lighting estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6908–6917, 2019. 2, 3, 4, 7
- [13] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 270–279, 2017. 5
- [14] Robin Green. Spherical harmonic lighting: The gritty details. In *Archives of the Game Developers Conference*, volume 56, page 4, 2003. 3
- [15] Yannick Hold-Geoffroy, Kalyan Sunkavalli, Sunil Hadap, Emiliano Gambaretto, and Jean-François Lalonde. Deep outdoor illumination estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7312–7321, 2017. 2, 6
- [16] Wenzel Jakob. Mitsuba renderer, 2010. 7
- [17] Antonios Karakottas, Nikolaos Zioulis, Stamatis Samaras, Dimitrios Ataloglou, Vasileios Gkitsas, Dimitrios Zarpalas, and Petros Daras. 360° surface regression with a hypersphere loss. In *2019 International Conference on 3D Vision (3DV)*, pages 258–268. IEEE, 2019. 1, 4
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [19] Jean-François Lalonde, Alexei A Efros, and Srinivasa G Narasimhan. Estimating natural illumination from a single outdoor image. In *2009 IEEE 12th International Conference on Computer Vision*, pages 183–190. IEEE, 2009. 2
- [20] Chloe LeGendre, Wan-Chun Ma, Graham Fyffe, John Flynn, Laurent Charbonnel, Jay Busch, and Paul Debevec. Deep-light: Learning illumination for unconstrained mobile mixed reality. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5918–5928, 2019. 2, 3, 4, 7
- [21] Marc Levoy, J Gerth, B Curless, and K Pull. The stanford 3d scanning repository. URL <http://www-graphics.stanford.edu/data/3dscanrep>, 5, 2005. 7
- [22] Zhengqin Li, Zexiang Xu, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Learning to reconstruct shape and spatially-varying reflectance from a single image. In *SIGGRAPH Asia 2018 Technical Papers*, page 269. ACM, 2018. 2
- [23] Stephen Lombardi and Ko Nishino. Reflectance and illumination recovery in the wild. *IEEE transactions on pattern analysis and machine intelligence*, 38(1):129–141, 2015. 2
- [24] Jorge Lopez-Moreno, Elena Garces, Sunil Hadap, Erik Reinhard, and Diego Gutierrez. Multiple light source estimation in a single image. In *Computer Graphics Forum*, volume 32, pages 170–182. Wiley Online Library, 2013. 2
- [25] Stephen R Marschner and Donald P Greenberg. Inverse lighting for photography. In *Color and Imaging Conference*, volume 1997, pages 262–265. Society for Imaging Science and Technology, 1997. 2
- [26] Abhimitra Meka, Christian Haene, Rohit Pandey, Michael Zollhöfer, Sean Fanello, Graham Fyffe, Adarsh Kowdle, Xueming Yu, Jay Busch, Jason Dourgarian, et al. Deep

reflectance fields: high-quality facial reflectance field inference from color gradient illumination. *ACM Transactions on Graphics (TOG)*, 38(4):77, 2019. [2](#)

- [27] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. [6](#)
- [28] Ravi Ramamoorthi and Pat Hanrahan. An efficient representation for irradiance environment maps. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 497–500. ACM, 2001. [3](#), [4](#)
- [29] Soumyadip Sengupta, Angjoo Kanazawa, Carlos D Castillo, and David W Jacobs. Sfsnet: Learning shape, reflectance and illuminance of faces in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6296–6305, 2018. [2](#)
- [30] Peter-Pike Sloan. Stupid spherical harmonics (sh) tricks. In *Game developers conference*, volume 9, page 42. Citeseer, 2008. [5](#)
- [31] Peter-Pike Sloan. Deringing spherical harmonics. In *SIGGRAPH Asia 2017 Technical Briefs*, page 11. ACM, 2017. [5](#)
- [32] Peter-Pike Sloan, Jan Kautz, and John Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *ACM Transactions on Graphics (TOG)*, volume 21, pages 527–536. ACM, 2002. [3](#)
- [33] Shuran Song and Thomas Funkhouser. Neural illumination: Lighting prediction for indoor environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6918–6926, 2019. [2](#), [3](#), [4](#), [6](#)
- [34] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1746–1754, 2017. [3](#), [7](#)
- [35] Tiancheng Sun, Jonathan T Barron, Yun-Ta Tsai, Zexiang Xu, Xueming Yu, Graham Fyffe, Christoph Rhemann, Jay Busch, Paul Debevec, and Ravi Ramamoorthi. Single image portrait relighting. *ACM Transactions on Graphics (TOG)*, 38(4):79, 2019. [2](#)
- [36] Henrique Weber, Donald Prévost, and Jean-François Lalonde. Learning to estimate indoor lighting from 3d objects. In *2018 International Conference on 3D Vision (3DV)*, pages 199–207. IEEE, 2018. [2](#)
- [37] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1851–1858, 2017. [1](#)
- [38] Nikolaos Zioulis, Antonis Karakottas, Dimitrios Zarpalas, Federico Alvarez, and Petros Daras. Spherical view synthe-

sis for self-supervised 360° depth estimation. In *2019 International Conference on 3D Vision (3DV)*, pages 690–699. IEEE, 2019. [5](#), [7](#)