

Fitting Spherical Gaussians to Dynamic HDRI Sequences

PASCAL CLAUSEN*, LI MA*, MINGMING HE*, AHMET LEVENT TAŞEL*, OLIVER PILARSKI*[†], and PAUL DEBEVEC, Netflix Eycline Studios, USA

We present a technique for fitting high dynamic range illumination (HDRI) sequences using anisotropic spherical Gaussians (ASGs) while preserving temporal consistency in the compressed HDRI maps. Our approach begins with an optimization network that iteratively minimizes a composite loss function, which includes both reconstruction and diffuse losses. This allows us to represent all-frequency signals with a small number of ASGs, optimizing their directions, sharpness, and intensity simultaneously for an individual HDRI. To extend this optimization into the temporal domain, we introduce a temporal consistency loss, ensuring a consistent approximation across the entire HDRI sequence.

ACM Reference Format:

Pascal Clausen, Li Ma, Mingming He, Ahmet Levent Taşel, Oliver Pilarski, and Paul Debevec. 2024. Fitting Spherical Gaussians to Dynamic HDRI Sequences. In *SIGGRAPH Asia 2024 Posters (SA Posters '24)*, December 3–6, 2024, Tokyo, Japan. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3681756.3698208>

1 INTRODUCTION

Environment lighting is commonly used to simulate distant illumination from a 360-degree scene. Compressed lighting representations reduce storage and computational costs in global illumination and inverse rendering. In this context, Ramamoorthi and Hanrahan [2001] use spherical harmonics (SH) for a linear approximation of low-frequency lighting. While this method is compact and efficient, it is limited to very low-frequency lighting. To address higher frequencies, Ng et al. [2003] introduce a nonlinear wavelet basis that retains only the largest terms, representing lighting across all frequencies. However, this method struggles with changing viewpoints and handles only diffuse surface interactions. Tsai and Shih [2006] propose using spherical radial basis functions (SRBFs) to model environments, capturing high-frequency signals through nonlinear optimization with scattered SRBFs.

Spherical Gaussians (SGs), like SH and SRBF, are rotation-invariant and compact, capable of representing both sharp and diffuse lights with relatively few parameters, enabling high compression. Xu et al. [2013] emphasize the user-friendly control of anisotropic SG (ASG) parameters and their efficient precomputation of light transfer functions using analytical approximations for integral, product, and

* Authors contributed equally to this work.

[†] Corresponding author

Authors' address: Pascal Clausen, pascal.clausen@scanlinevfx.com; Li Ma, lmaag@connect.ust.hk; Mingming He, hmm.lillian@gmail.com; Ahmet Levent Taşel, leventtasel@gmail.com; Oliver Pilarski, oliver.pilarski@scanlinevfx.com; Paul Debevec, debevec@gmail.com, Netflix Eycline Studios, USA.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SA Posters '24, December 3–6, 2024, Tokyo, Japan

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1138-1/24/12

<https://doi.org/10.1145/3681756.3698208>




Fig. 1. Our method approximates an HDRI sequence using a temporally consistent mixture of anisotropic spherical Gaussians (ASGs). Below are three example frames showing our fitting results and the rendering of balls with varying roughness under the corresponding HDRI.

convolution operations. Yet, we find that the L2 loss used in these optimizations ([Tsai and Shih 2006; Xu et al. 2013]) performs poorly with low-frequency lighting. We propose a new loss function that combines L1 reconstruction loss with a diffuse loss to better preserve all frequencies. While the rotation invariance of ASGs provides stable approximations for rotated HDRIs, frame-by-frame approximation of HDRI sequences, such as those captured over a short period ([Stumpfel et al. 2004]), lacks temporal stability. To address this issue, we introduce a temporally consistent approximation using a temporal consistency term, ensuring stable and consistent HDRI sequences over time.

In summary, this work introduces a technique to approximate a time sequence of high dynamic range natural illumination using a small number of ASGs, with the following technical contributions:

- (1) An optimization algorithm featuring a novel loss function that combines diffuse loss and L1 reconstruction loss, effectively capturing both low-frequency and high-frequency information in HDRIs.
- (2) An extension of the optimization to the temporal domain, enabling the approximation of temporally consistent environment lighting.

2 DESIGN AND IMPLEMENTATION

To approximate an HDRI using a mixture of ASGs, we represent each frame of the HDRI sequence with a set of ASGs, keeping the number of ASGs constant but varying their parameters. For each ASG, we use the same definition as the *Bingham distribution* in [Xu et al. 2013]. Each ASG $G_i(\mathbf{d})$ is defined by the formula:

$$G_i(\mathbf{d}; [\mu, \lambda], [\mathbf{u}, \mathbf{v}], \mathbf{c}) = c e^{-\mu(\mathbf{d} \cdot \mathbf{u}) - \lambda(\mathbf{d} \cdot \mathbf{v})}, \quad (1)$$

where \mathbf{d} is the query direction vector, \mathbf{u}, \mathbf{v} being the tangent and bi-tangent axes, respectively. $[\mu, \lambda]$ are the bandwidths along the two axes, indicating the sharpnesses in these directions. c denotes the peak intensity of the ASG. We parameterize $[\mathbf{u}, \mathbf{v}]$ by $[\mathbf{u}, \mathbf{n}]$, where \mathbf{n} is the direction of the gaussian lobe and \mathbf{v} is derived from those directions. Each frame of the HDRI sequence is represented by a mixture of ASGs: $I_{pred}(\mathbf{d}) = \sum_i G_i(\mathbf{d})$, where I_{pred} denotes

the predicted HDRI. To optimize temporally consistent ASGs, our algorithm involves the following steps:

- For the initial frame of the HDRI time sequence, we optimize the ASG parameters (directions, sharpness, intensity) using an Adam optimizer for 24'000 epochs.
- For each consecutive frame, we initialize the ASG parameters with those from the previous frame and then optimize them using an Adam optimizer with an additional temporal regularization loss for 6'000 epochs.

At each optimization iteration, we densely sample 256×512 pixels from the HDRI with equirectangular projection, and optimize the Gaussian parameters $\mathbf{g}_i = (\mu_i, \lambda_i, \mathbf{u}_i, \mathbf{n}_i, c_i)$ using an analysis-by-synthesis approach. We minimize a composite loss function $L = w(\alpha L_R + \beta L_D) + \gamma L_T$, where w is the solid angle of each pixel. We empirically set $\alpha = 1.0$, $\beta = 1.0$ and $\gamma = 0.5$ in our experiments. L_R , L_D , L_T are the reconstruction loss, diffuse loss, and temporal consistency loss given by:

$$L_R = \|I_{pred} - I_{gt}\|_1, L_D = \|D_{pred} - D_{gt}\|_1,$$

$$L_T = \sum_i \|(\mathbf{g}_i^t - \mathbf{g}_i^{t-1}) / \max\{\mathbf{g}_i^{t-1}\}\|_2,$$

where I_{pred} and I_{gt} are the predicted and ground truth HDRI intensities at the sampled directions. Our experiments show that using L1 loss effectively fits sharp edges. D_{pred} and D_{gt} denote the predicted and ground truth intensities of a diffused HDRI version at the sampled directions, which is a Lambertian light map approximated using spherical harmonics (SH) of degree 3, with SH coefficients obtained through numerical integration. The diffuse loss is crucial for optimizing ASGs to maintain correct total energy. \mathbf{g}_i^t and \mathbf{g}_i^{t-1} are parameters of i -th ASG for the current and previous frame.

3 RESULTS

We present the fitting results of HDRI along with renderings of balls using different numbers of ASGs, compared to their ground truths, in Fig. 2. It can be observed that beyond a small number of ASGs (e.g. 15), the differences in the rendered balls at most levels of roughness become negligible when using the compressed HDRI for environment lighting, compared to using the ground truth HDRI. This indicates that our method effectively reconstructs the HDRI with a relatively low number of ASGs. Moreover, the fidelity of the rendering results across different roughness levels indicates that the compressed HDRI is well preserved for different frequencies. Our method is also applicable to fitting with isotropic spherical Gaussians. Please refer to the accompanying video for more results.

To validate our composite loss design, we first compare different loss functions for single HDRI fitting, as shown in Fig. 3. The experiments demonstrate that L1 loss performs best for capturing both low- and high-frequency lighting signals in the HDRI, as evidenced by the diffuse and specular shading effects in the renderings. In contrast, using L2 loss, as in [Tsai and Shih 2006] or [Xu et al. 2013], even with HDRI intensity preprocessing through functions like log or sRGB, does not yield good results, likely due to the extreme dynamic range of the HDRI. Additionally, without the diffuse loss L_D , it leads to inaccurate recovery of low-frequency lighting and also negatively impacts the high-frequency information, as diffuse loss is important for preserving the global energy of the HDRI. Finally,

we demonstrate the necessity of using the temporal consistency loss L_T to maintain temporal stability, as shown in Fig. 4. It displays time variation vertically for a selected row from the latlong map horizontally. Without this loss, strong flickering occurs, which is even more noticeable in the video.




Fig. 2. The comparison of the ground truth HDRI and compression results with different numbers of ASGs, shown through latlongs and renderings.




Fig. 3. The comparison of different loss functions used for HDRI fitting with 15 ASGs. L1 refers to the L1 loss function used in the reconstruction loss; L2 indicates that the L1 loss is replaced by L2 loss; L_D is the diffuse loss.




Fig. 4. The comparison between with and without temporal consistency loss L_T using 15 ASGs. The selected row from the HDRI sequence are stacked along the time axis to visualize temporal consistency.

4 CONCLUSION AND FUTURE WORK

We propose an optimization method for temporally consistent dynamic HDRI compression preserving all frequencies using ASGs. However, it struggles with fine details in mirror reflections and cannot accurately render sharp-angled lights (e.g. rectangular lights) due to the mixing of spherical Gaussians with higher components.

REFERENCES

- Ren Ng, Ravi Ramamoorthi, and Pat Hanrahan. 2003. All-frequency shadows using non-linear wavelet lighting approximation. *ACM Trans. Graph.* 22, 3 (2003), 376–381. <https://doi.org/10.1145/882262.882280>
- Ravi Ramamoorthi and Pat Hanrahan. 2001. An efficient representation for irradiance environment maps. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 2001, Los Angeles, California, USA, August 12-17, 2001*, Lynn P. Coeck (Ed.). ACM, 497–500. <https://doi.org/10.1145/383259.383317>
- Jessi Stumpfel, Andrew Jones, Andreas Wenger, Chris Tchou, Tim Hawkins, and Paul E. Debevec. 2004. Direct HDR capture of the sun and sky. In *International Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 2004, Los Angeles, California, USA, August 8-12, 2004, Posters*, Ronen Barzel (Ed.). ACM, 50. <https://doi.org/10.1145/1186415.1186473>
- Yu-Ting Tsai and Zen-Chung Shih. 2006. All-frequency precomputed radiance transfer using spherical radial basis functions and clustered tensor approximation. *ACM Trans. Graph.* 25, 3 (2006), 967–976. <https://doi.org/10.1145/1141911.1141981>
- Kun Xu, Wei-Lun Sun, Zhao Dong, Dan-Yong Zhao, Run-Dong Wu, and Shi-Min Hu. 2013. Anisotropic spherical Gaussians. *ACM Trans. Graph.* 32, 6 (2013), 209:1–209:11. <https://doi.org/10.1145/2508363.2508386>