

# Deep Illumination–Guided Light Probe Placement

A. Tarasidis<sup>1</sup> and A. A. Vasilakis<sup>2</sup>  and I. Fudos<sup>1</sup> 

<sup>1</sup>Department of Computer Science and Engineering, University of Ioannina, Ioannina, Greece

<sup>2</sup>Department of Informatics, Athens University of Economics and Business, Greece

## Abstract

*This work proposes an automated learning-based strategy for computing light probe layouts efficiently under varied illumination conditions. A neural network model estimates the relative contribution of candidate probes, enabling the rapid construction of a compact configuration that maintains the scene’s indirect lighting distribution. Evaluations on complex environments indicate that the method achieves substantial speedups over conventional placement methods without compromising illumination fidelity.*

## CCS Concepts

• **Computing methodologies** → *Neural networks; Rendering;*

## 1. Introduction

Real-time radiance caching is highly popularized in the games industry due to its relative simplicity and efficiency to incorporate indirect illumination phenomena. Nowadays, most industrial rendering engines have integrated light probes (LP) creation to achieve run-time global illumination on dynamic objects [Uni25]. The most common automated strategy consists of deploying a dense, regular grid of LP, followed by local adjustments to better align LPs with regions of high color contrast or geometric detail. However, this approach often fails to preserve global lighting behavior when resources are limited, as uniform sampling tends to undersample fine-scale geometric features while overspending LPs in less relevant areas. Increasing grid resolution can mitigate some reconstruction artifacts, but at the cost of higher memory usage and additional, time-consuming tuning. Subsequent efforts investigated adaptive or heuristic-driven placement strategies [WKKN19, VVP21]. However, many of these techniques still depend on manual intervention or extensive parameter tuning, which limits their robustness and scalability in production pipelines. In this work, we propose a neural network–driven framework for accelerating LP placement in complex scenes. Our method learns to estimate LP relevance directly from scene and lighting features, enabling the rapid construction of compact LP layouts that preserve indirect illumination fidelity under constrained memory budgets.

## 2. Method Overview

Our method accelerates LP placement in complex scenes by predicting the relative importance of candidate LPs and selecting only those that have the greatest impact on the scene’s indirect illumination. We employ a PointNet-based neural network [QSMG17], referred in our case as the Light-Probe Neural Network (LPNN)

to infer an importance score for each position in a dense LP grid. These scores are computed from a set of geometric and photometric features that capture illumination variation, and visibility complexity. The predicted importance score quantifies how significantly a LP contributes to the overall indirect illumination observed by evaluation light points (ELPs) uniformly scattered through the scene. Based on this score, we select the most impactful LPs (subset of ELP) using a percentile-based threshold, allowing users to directly control the desired sparsity of the final LP layout.

**Feature & Label Collection.** We adopt a grid-based sampling approach to systematically evaluate each voxel in the scene. From each of the  $N$  ELP, a set of *features* is extracted to capture both photometric and visibility characteristics of the surrounding scene. Specifically, this feature vector consists of 30 floating-point values: *Spherical Harmonic L2* (24), *Luminance variance* (3), *Color variance* (3), *Normal variance* (1), *Occlusion factor* (1). For the *labels* required by our supervised-learning framework, any existing LP layout can be used as ground truth. We employ a fully agnostic labeling algorithm that identifies the presence of a LP in the vicinity of ELPs. A LP near an ELP is considered to represent a location of high importance. In this work, we adopt the ground-truth light probe layouts produced by [VVP21].

**Light Probe Neural Network.** The LPNN model begins with three sequential Convolution1D layers, each followed by Batch Normalization, expanding the input feature vector from length 30 to 256 dimensions. These layers progressively increase channel size from 64 to 128 to 256 using ReLU activations. To incorporate global context, the network applies both Global Max Pooling 1D and Global Average Pooling 1D, concatenating these with local point features to produce a 512-dimensional representation. A custom tiling function then replicates this global context across local neighborhoods, ensuring each point’s importance is influenced primar-

ily by nearby lighting data, while still retaining some awareness of the overall scene due to pooling before tiling. After context tiling, the model applies an additional Convolution1D sequence where Dropout layers replace Batch Normalization to reduce overfitting by randomly disabling 30% of neurons during training. This sequence compresses the high-dimensional features into a single importance score in the  $[0, 1]$  range via a final Convolution1D layer with a Sigmoid activation. The model is trained for 50 epochs on 3,485 feature vectors gathered from multiple scenes with varied layouts and label settings, enabling it to generalize across different lighting scenarios. Importantly, the input features are scene-agnostic, containing no explicit geometric or lighting source information, allowing the model to adapt to unseen environments without overfitting, while simplifying the dataset expansion process, even using the same scene with a variety of initial grid layouts.

**Light Probe Selection.** After inference is completed, the initial ELP group is populated based on the model’s predicted importance scores, but these scores often exhibit clustering around certain values, making direct thresholding less effective. To address this, the values are first remapped to a normalized scale  $[0 - 1]$ , ensuring the full range is utilized and reducing the impact of clustering. The placement system uses a percentile-based thresholding method  $R_t$  that enables the user to define the proportion of total probes to be placed, prioritizing those with the highest scores.

### 3. Evaluation

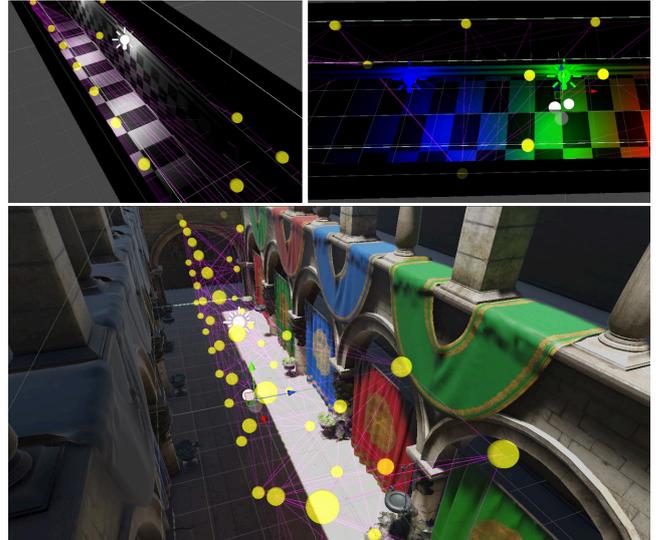
This section presents the experimental evaluation of our approach, with an emphasis on quantitative and qualitative assessment. We developed a cross-platform prototype of our approach as a component within the Unity game engine. This integration allows us to leverage Unity’s light probe baking pipeline, as well as its tetrahedral tessellation system for probe graph construction, connectivity, and interpolation [Uni25]. A preliminary version of the prototype is available at <https://github.com/AndreasTar/LPNN>.

**Performance.** Table 1 shows that LPNN performs orders of magnitude faster than [VVP21], with placement times remaining consistently within a few seconds regardless of the number of EVL or scene complexity. In contrast, [VVP21] exhibited significantly longer and more variable execution times depending on grid and decimation settings. However, the LPNN method can occasionally place probes in non-critical areas, leading to oversampling and requiring minimal manual adjustment.

**Quality.** A high-quality LP layout strikes a balance between lighting fidelity and resource efficiency. LPs should be concentrated

**Table 1:** Execution comparison of [VVP21] (left) and our method (right) for different scenes and settings; ELP count ( $N$ ), reduced LPs ( $M$ ), evaluation points ( $EP$ ), reduction  $R_t$ .

Scene	Timing	$N$ ( $M$ )	$EP$	$R_t$
Sponza	22.4s / <b>5.3ms</b>	105 (34)	128	0.4
	600.3s / <b>17.8ms</b>	240 (54)	256	0.9
Corridor	161.1s / <b>10.9ms</b>	180 (120)	256	0.5
	477.6s / <b>15.7ms</b>	243 (147)	256	0.6



**Figure 1:** LP layouts for the Corridor (top) and Sponza (bottom).

in areas where illumination changes rapidly, such as near shadow boundaries, occluding geometry, or regions with strong color bleeding, as these are the points most susceptible to interpolation error in irradiance data. Conversely, uniform lighting regions, such as open, evenly lit spaces, can be sparsely populated without noticeable quality degradation. By adapting probe density to local lighting variance and geometric complexity, it is possible to capture global illumination effects with fewer probes while still approximating a path-traced ground truth. We tested our approach on scenes of different radiometric complexity and observed that we obtained sufficient light probe count reduction under diverse lighting conditions and constraints, as illustrated in the Figure 1.

### 4. Future Work

Future directions could focus on refining the neural network architecture, particularly by incorporating features sensitive to probe edge boundaries to improve placement accuracy and minimize under/over-sampling. Enhancements in network design, paired with more diverse and representative training datasets, may improve generalization while reducing computational cost.

### References

- [QSMG17] QI C. R., SU H., MO K., GUIBAS L. J.: PointNet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE CVPR* (2017). doi:10.1109/CVPR.2017.16. 1
- [Uni25] UNITY: Adaptive probe volumes, 2025. 1, 2
- [VVP21] VARDIS K., VASILAKIS A. A., PAPAIOANNOU G.: Illumination-driven light probe placement. In *Eurographics 2021 - Posters* (2021), The Eurographics Association. doi:10.2312/egp.20211026. 1, 2
- [WKKN19] WANG Y., KHIAT S., KRY P. G., NOWROUZEZAHRAI D.: Fast non-uniform radiance probe placement and tracing. In *Proceedings of the ACM SIGGRAPH Symposium on I3D* (NY, USA, 2019), ACM. doi:10.1145/3306131.3317024. 1