EUROGRAPHICS AACHEN 2026

DEEP ILLUMINATION-GUIDED LIGHT PROBE PLACEMENT

Andreas Tarasidis[1], Andreas A. Vasilakis[2] and Ioannis Fudos[1]
[1]University of Ioannina, Greece
[2]Athens University of Economics and Business, Greece

PROBLEM

Real-time 3D applications require efficient **global illumination** (GI) to achieve believable lighting under strict performance constraints. **Light Probes (LPs)** provide a practical solution for approximating indirect lighting, but their effectiveness strongly depends on their spatial distribution [Greger98]. Naïve approaches, such as **uniform probe grids**, waste resources in low-importance regions while failing to capture complex lighting variations. More advanced methods improve placement using **heuristics or adaptive tuning**, but they often lack scalability and robustness across diverse scenes.

GOAL: The core challenge is to determine a **compact set of probe positions** that **preserves indirect lighting** quality while minimizing memory usage and maintaining real-time performance.

RELATED WORK

This is traditionally addressed through an **iterative trial-and-error** process in which experienced artists **manually place** and adjust LP locations until an acceptable visual outcome is achieved. Moreover, any subsequent **modification to scene geometry** or lighting invalidates the existing layout, requiring the entire placement process to be **repeated**.

The most common automated strategy consists of deploying a **dense, regular grid** of LP. However, uniform sampling tends to under-sample fine-scale geometric features while overspending LPs in less relevant areas [Tatarchuk05]. Prior work has explored **adaptive** and **heuristic-based** placement strategies, leveraging geometric complexity [Wang19], visibility [Majercik21], or radiance variance [Vardis21] to better place probes. While they improve efficiency, they typically rely on iterative optimization, or parameter tuning, **limiting their robustness** across diverse scenes.

OVERVIEW

We propose a **neural network–driven approach** for **efficient LP placement** by retaining only those that contribute most to the indirect illumination:

1. **Dense Sampling:** Generate a uniform grid of Evaluation LP (ELP)
2. **Feature Extraction**: Compute geometric and photometric features & Capture illumination variation and visibility complexity
3. **Importance Prediction:** Our Light-Probe Neural Network outputs an importance score per ELP
4. **Probe Selection:** Select top-ranked LP using a percentile threshold value

METHODOLOGY

1. **Uniform ELP Sampling Strategy**
- Structured grid layout ensuring **complete scene coverage**
- **Uniform baseline** for subsequent reduction

2. **Feature Generation**
- **30 dimensional vector per ELP**, encodes **lighting & geometric** context:
  - L2 Spherical Harmonics
  - Luminance, Color, Normal, Occlusion Variance
- **Supervised learning:**
  - Detect LP presence near each ELP
  - Ground truth based on [Vardis21]
- Scene agnostic, **no dependencies** on geometry/lighting

3. **Light-Probe Neural Network [PointNet2017]**
- **Input:** 30D feature vector per probe candidate
- **Feature extraction** via Conv1D stack:
  - Channels: 64 → 128 → 256
  - Activation: ReLU, each layer followed by Batch Normalization
- **Output:** 256D local feature representation
- **Global Context Integration** via Global Max &Average Pooling
- **Prediction Head** to reduce overfitting
  - Additional Conv1D layers to dropout 30%
- **Final layer** via Conv1D + Sigmoid activation
  - Outputs importance score

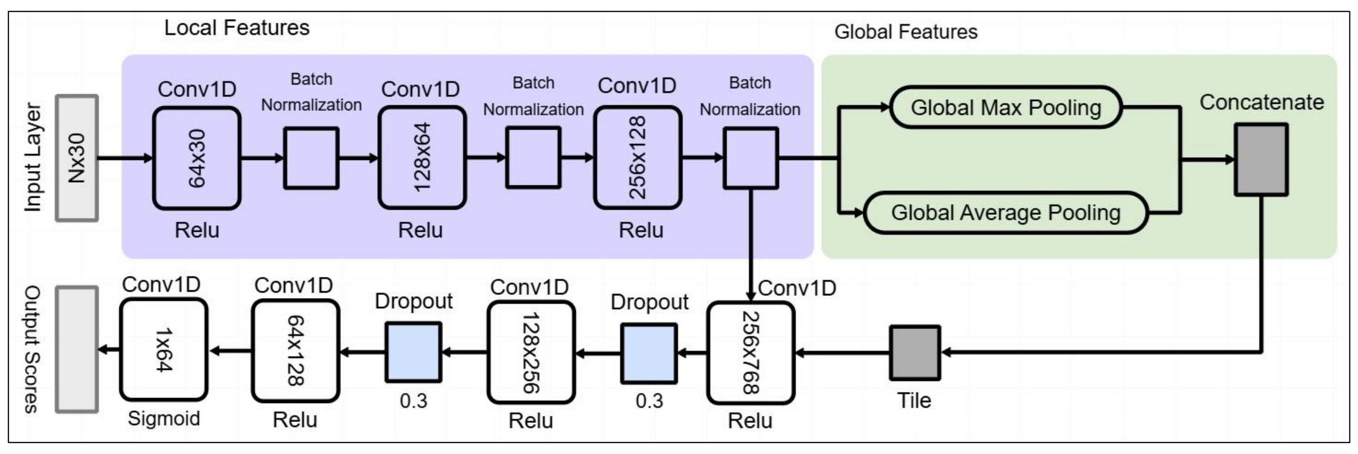

4. **Light Probe Selection**
- Initial ELP set is populated using **predicted importance scores**
- **Normalization** to [0-1] is performed to ensure full range utilization
- User select top-performing subset by **important score thresholding**
- Predictable and **tuneable placement** behavior

RESULTS

Evaluation through both **quantitative** and **qualitative** analysis, focusing on the **performance**, **memory** and **quality** of the resulting LP configurations.

Performance:
- Our method achieves **orders-of-magnitude speedup** over [Vardis21]
- LP placement in a **few seconds**, independent of ELPs & scene complexity

Memory:
- Memory usage scales **linearly** with number of ELPs
- Storage remains **minimal**: ≈1 KB per 100 ELPs

Quality:
- **Preserve indirect lighting** detail with reduced probe count
- **Dense placement** in regions with rapid illumination changes, shadow boundaries, occlusions, color bleeding
- **Sparse sampling** in uniformly lit areas & open, low-variance regions

Implementation Details: Cross-platform prototype within **Unity 6**, that allows us to leverage LP baking pipeline and its tetrahedral tessellation system for LP graph construction, connectivity, and interpolation [Uni25].

Experimental Setup:
- **CPU:** Intel i7-9750H CPU, **GPU:** NVIDIA 2060M GPU, **OS:** Windows 10

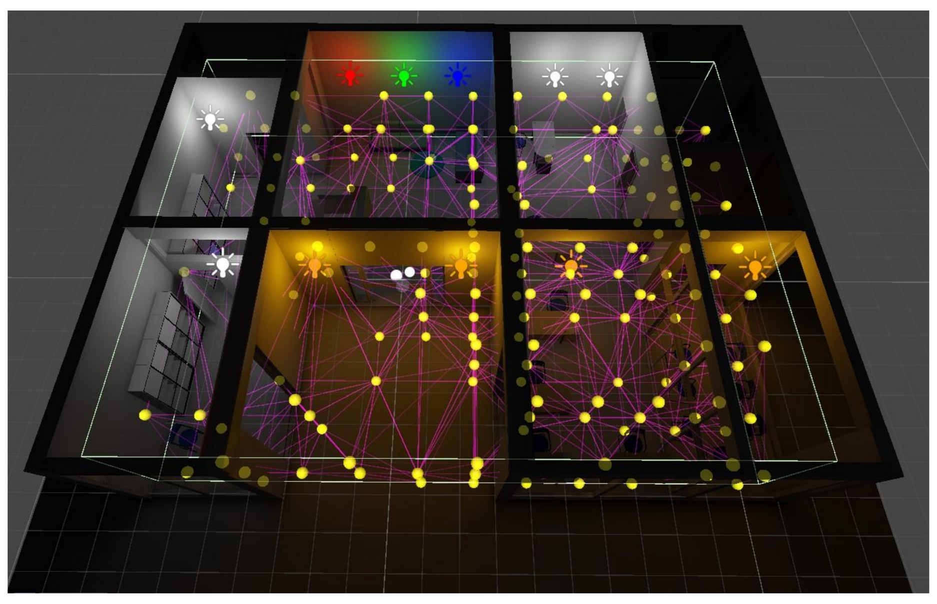

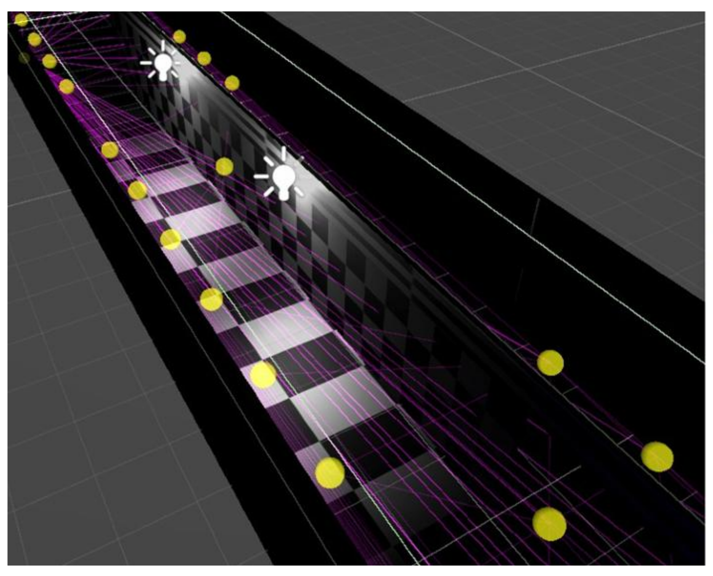

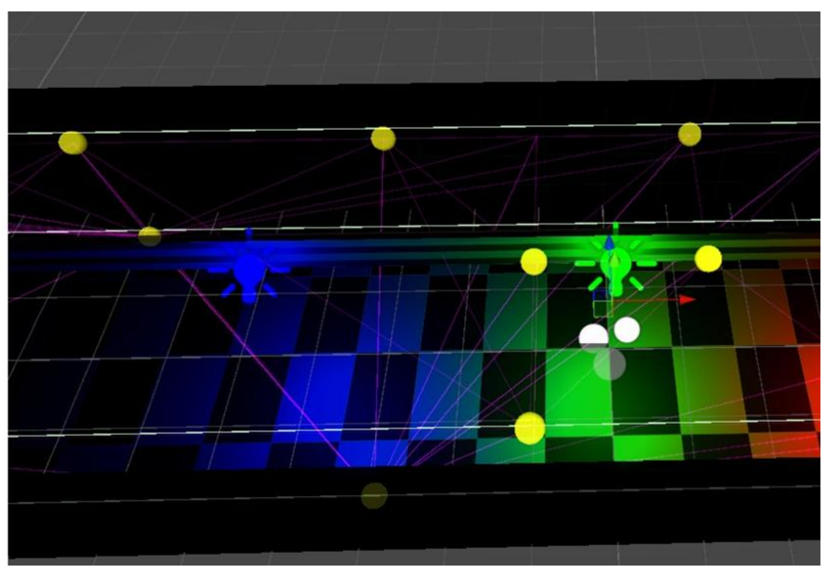

AFFILIATIONS

ΟΠΑ AUEB

AUEB COMPUTER GRAPHICS GROUP

CGRG Computer Graphics Research Group

REFERENCES

[Greger98] **The irradiance volume**, CG&A, 1998.
[Tatarchuk05] **Irradiance volumes for games**, GDC 2005.
[Qi17] **PointNet: Deep learning on point sets for 3d classification and segmentation**, CVPR 2017.
[Wang19] **Fast non-uniform radiance probe placement and tracing**, I3D 2019.
[Vardis21] **Illumination-driven light probe placement**, EG 2021 (posters section).
[Majercik21] **Scaling Probe-Based Real-Time Dynamic Global Illumination for Production**, JCGT, 2021.

LINKS

Source code:
- *https://github.com/AndreasTar/LPNN*

Computer Graphics Groups:
- *https://cgrg.eu/*
- *https://graphics.cs.aueb.gr/*