



Neural importance sampling of many lights

Bastien DOIGNIES - GDL - LIRIS - 01/08/2025


Rappel path tracing



Rappel path tracing




Rappel path tracing



Choosing light source - Uniform sampling

First idea : uniform sampling of lights.




What if the point is not lit by the chosen light?

1. Because the contribution is not relevant
2. Because there is an occluder

RIS [Talbot 2005] / ReSTIR [Bitterli 2020]


Idea: gather information from multiple lights

1. Select some “random” candidates
 2. Evaluate their importance -> weights
 3. Resample theses candidates with weights
-> One sample becomes better
-
4. Reuse information across neighborhood and time, and whatever you wish



Proportional to BRDF / cosines / ...
/!\ But not visibility

Framework





Estimate light cluster probability with network and combine it with other methods.

Sample light within cluster

Use light for path tracing

Neural network : inputs (and encoding)

TLDR: position, brdf value and normal



Neural network : loss


TLDR: KL divergence between the network and the “received” spectral radiance i.e. the cluster importance

The gradient is:

$$\langle \nabla_{\theta} D_{\text{KL}}(q, p_{\theta}) \rangle = -\frac{1}{N} \sum_{j=1}^N \left[\frac{F(\mathbf{x}, \mathbf{Z}_j, \omega_o)}{p(\mathbf{Z}_j | Y_j) p(Y_j | C_j) p_{\theta}(C_j)} \nabla_{\theta} \log p_{\theta}(C_j) \right].$$

Neural network : training

1. Online training (can be cut after some spp)
2. Only estimates residuals: at initialization, the network is dreadfully bad at choosing light sources



Results - FLIP metric






Fig. 1. The FLIP pipeline.

Results - equal time (5s)



Results - Equal samples (128spp)



Results

Table 1. Equal-time comparisons measured in FLIP [Andersson et al. 2020] with number of samples per pixel for eight representative scenes. We compare our approach against ATS [Conty Estevez and Kulla 2018], SLCRT [Lin and Yuksel 2020], ReSTIR [Bitterli et al. 2020], and VARL [Wang et al. 2021]. We color code the **first**, **second**, and **third** lowest numbers.

	ATS		SLCRT		ReSTIR		VARL		Ours	
BATHROOM	0.1659	212	0.1413	220	0.1744	135	0.1346	139	0.0977	175
BEDROOM	0.3129	192	0.2442	196	0.2796	123	0.1960	123	0.0910	157
BISTRO	0.1548	163	0.2656	165	0.2837	74	0.2281	60	0.1126	168
LIVING ROOM	0.1247	194	0.0977	194	0.1314	131	0.1009	131	0.0671	146
SAN MIGUEL	0.2253	159	0.3015	158	0.2817	110	0.2472	107	0.1713	164
STAIRCASE	0.0763	198	0.0817	203	0.1059	129	0.0946	138	0.0529	160
STAIRCASE2	0.1020	233	0.0925	233	0.1132	159	0.0854	149	0.0646	191
ZERO DAY	0.1647	222	0.1455	223	0.1682	159	0.1338	161	0.1117	201

Table 2. Equal-sample comparisons measured in FLIP [Andersson et al. 2020] with runtime in seconds. All the scenes are rendered using each approach for 128 samples per pixel.

	ATS		SLCRT		ReSTIR		VARL		Ours	
BATHROOM	0.2059	2.8	0.1791	2.9	0.1780	4.7	0.1399	4.3	0.1090	3.9
BEDROOM	0.3780	3.2	0.2962	3.2	0.2757	5.2	0.1913	4.8	0.0994	4.4
BISTRO	0.1704	15.9	0.2944	15.3	0.2316	36.8	0.1658	44.3	0.1180	19.6
LIVING ROOM	0.1488	3.2	0.1162	3.3	0.1387	5.2	0.1023	4.6	0.0689	4.8
SAN MIGUEL	0.2454	7.9	0.3233	8.1	0.2675	12.4	0.2308	11.7	0.1842	8.6
STAIRCASE	0.0923	3.3	0.1002	3.2	0.1060	4.9	0.0969	4.7	0.0553	4.5
STAIRCASE2	0.1246	2.7	0.1128	2.8	0.1206	4.0	0.0913	4.1	0.0721	3.8
ZERO DAY	0.1929	5.1	0.1699	5.0	0.1771	8.1	0.1465	7.0	0.1224	7.0

FIN