

Real-time Rendering with a Neural Irradiance Volume

Arno Coomans^{*1}  Giacomo Nazzaro^{*1}  Edoardo A. Dominici¹  Christian Döring² 
Floor Verhoeven¹  Konstantinos Vardis¹  Markus Steinberger^{3,4} 

Huawei Technologies, Switzerland¹, Germany², Austria³
Graz University of Technology, Austria⁴

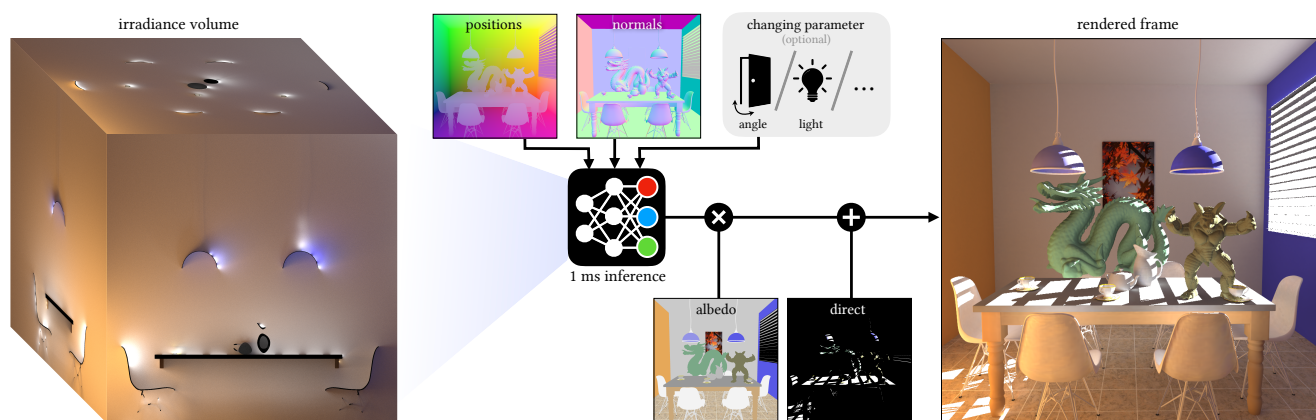


Figure 1: Overview of rendering dynamic objects in a static scene using our Neural Irradiance Volume. Left: The irradiance field of the scene is visualized through three orthogonal slices of its volume. Center: A lightweight neural model compresses the 5D irradiance field for the whole scene's volume. By querying the model with position and normal buffers as input, indirect diffuse illumination is computed at real-time rates. Right: The model's output is efficiently utilized to compute high-quality diffuse global illumination, not only on static surfaces but also on novel dynamic objects—such as the Dragon and Armadillo above—which are effectively immersed in the neural irradiance volume. Optionally, the model can be trained with changing scene parameters, obtaining non-static irradiance fields.

Abstract

Rendering diffuse global illumination in real-time is often approximated by pre-computing and storing irradiance in a 3D grid of probes. As long as most of the scene remains static, probes approximate irradiance for all surfaces immersed in the irradiance volume, including novel dynamic objects. This approach, however, suffers from aliasing artifacts and high memory consumption. We propose Neural Irradiance Volume (NIV), a neural-based technique that allows accurate real-time rendering of diffuse global illumination via a compact pre-computed model, overcoming the limitations of traditional probe-based methods, such as the expensive memory footprint, aliasing artifacts, and scene-specific heuristics. The key insight is that neural compression creates an adaptive and amortized representation of irradiance, circumventing the cubic scaling of grid-based methods. Our superior memory-scaling improves quality by at least 10x at the same memory budget, and enables a straightforward representation of higher-dimensional irradiance fields, allowing rendering of time-varying or dynamic effects without requiring additional computation at runtime. Unlike other neural rendering techniques, our method works within strict real-time constraints, providing fast inference (around 1 ms per frame on consumer GPUs at full HD resolution), reduced memory usage (1–5 MB for medium-sized scenes), and only requires a G-buffer as input, without expensive ray tracing or denoising.

Keywords: real-time, rendering, global illumination, pre-computation.

CCS Concepts

• **Computing methodologies** → **Ray tracing; Neural networks;**

* Both authors contributed equally to the paper.




Figure 2: In many scenes, indirect lighting contributes the majority of visible radiance. By caching indirect lighting using an irradiance volume—such as our neural irradiance volume—the significant cost of Monte Carlo integration at runtime can be avoided.

1. Introduction

Rendering moving objects within static scenes is a common requirement in real-time applications, where characters and animated entities interact with predefined environments. However, computing global illumination at high frame rates remains challenging due to the non-local and computationally expensive nature of indirect lighting.

A common solution in practice is to “bake” the illumination of the static scene into a grid of light probes, forming an irradiance volume that can be queried efficiently at runtime to shade unseen objects. Probe-based irradiance volumes have long constituted the de facto standard for global illumination in video games, valued for their simplicity and runtime efficiency. While not physically accurate, their rendered effects are visually convincing for all practical real-time applications. However, probe-based methods are not without limitations: placement and density must typically be determined using scene-specific heuristics, and interpolation between probes can introduce noticeable artifacts. Moreover, because the representation is inherently low-frequency, probes struggle to capture fine-scale details such as contact shadows or subtle irradiance variations. To achieve high-frequency lighting effects, the density of the probe grid needs to be significantly increased, which quickly becomes intractable due to their cubic scaling.

Beyond classical probes, recent neural methods have also demonstrated how neural models can accelerate rendering by caching radiative quantities on surfaces [HCZ21, RWG*13, CDD*24, SDR*24, DPD22]. While these methods achieve impressive fidelity, they are generally unsuitable for real-time applications involving dynamic objects unseen during training, as pre-computation only accounts for known scene surfaces. Similarly, neural caches have been employed to amortize real-time path tracing costs through online learning [MRNK21], but the reliance on expensive steps such as denoising and ray tracing makes them impractical under the strict runtime budgets of real-time applications. However, as support for neural operations becomes increasingly widespread and performant, reframing real-time rendering problems with

neural solutions allows them to benefit from the acceleration of tensor hardware performance. In this work, we propose a Neural Irradiance Volume (NIV), which evolves the concept of irradiance volumes by leveraging recent neural advances to deliver high-quality, real-time indirect diffuse illumination at low memory, for both dynamic objects and static surfaces. NIV is inspired by the foundational work on irradiance volumes [GSHG98], adopting the same assumptions common in practical applications [BB*17, SSS*20]. Namely, the scene is considered mainly *static* at runtime, while dynamic elements—such as characters or interactive entities—are relatively small and do not significantly impact global light propagation.

Our method retains many of the runtime advantages of classical probes while enabling the superior shading quality and memory use of neural methods. Unlike probe-based methods, our approach is heuristic-free, relying instead on backpropagation to directly optimize for reconstruction quality. We tackle the cubic scaling of irradiance volumes by encoding the irradiance field in a compact neural model, achieving stronger compression without sacrificing high-frequency details. Unlike neural methods, we compute a pre-integrated and volumetric quantity, which permits noise-free shading of dynamic objects, without the use of ray tracing or denoising.

Our contributions are the following:

- A precomputation scheme that bakes path tracing quality irradiance without requiring per-scene heuristics, with the ability to shade unseen moving objects.
- An order of magnitude lower rendering cost compared to state-of-the-art neural rendering methods.
- A compact representation for irradiance fields that, at a fixed memory budget, improves quality by an order of magnitude over competing real-time volumetric approaches.
- The ability to extend our method to non-static irradiance fields, which is a common application, e.g. a time-of-day cycle.

2. Related Work

Traditional Precomputation. Computing a noise-free solution to the rendering equation [Kaj86] of arbitrary scenes remains infeasible in real-time. Consequently, rendering research has long relied on preprocessing techniques, or “baking”, to eliminate parts of this computation at runtime. Since direct illumination is comparatively inexpensive to compute, many methods focus on precomputing the more complex, indirect component of (diffuse) light transport.

Our work draws significant inspiration from irradiance probe volumes [GSHG98], sharing the goal of volumetrically precomputing irradiance to enable efficient and stable shading of moving objects at runtime. A multitude of follow-up work has improved probe-based irradiance caching, focusing on the challenges of probe placement and interpolation, to mitigate light-leaking artifacts and reconstruction errors. Solutions include heuristics like the use of geometric terms [ST15], normal-offsetting [Hoo16], non-uniform probe placement [WKK19, ZWZ*25], and compression [VPG14, ZLZ*25]. Enhancements such as visibility information [IS17, MMNL17] and precomputed radiance transfer [SKS02] have further refined shading accuracy through the use of occlusion information, while extensions to glossy materials [RLP*20] have




Figure 3: Diffuse global illumination is inexpensive to render starting from indirect irradiance E . However, the indirect irradiance signal itself is expensive to compute, making it the most valuable component to cache. Multiplying the latter by surface albedo ρ results in indirect illumination, and adding direct illumination D and emitted radiance L_e results in diffuse global illumination.

expanded applicability. Unlike probe-based techniques, our method represents irradiance continuously, forgoing the problems of probe interpolation and placement.

Precomputation using Neural Networks. Neural networks have been extensively studied for approximating radiance, with applications ranging from static scenes [HCZ21] to those involving dynamic parameters such as scene illumination [RWG*13, RBRD22], various scene configurations [SDR*24, DPD22] or manifolds within a set of scene parameters [CDD*24]. In terms of architecture, more advanced neural representations have recently been explored, such as those based on a transformer architecture [RHP*24, ZDP*25]. A key advantage of all the aforementioned neural methods is their ability to learn representations implicitly, eliminating the need for problem-specific heuristics required by traditional precomputation approaches. Neural techniques typically deliver higher quality than classical methods, though they come with longer precomputation times and higher inference costs. More recently, attempts have been made to combine neural and traditional precomputation [ISSS24], by equipping each probe with a small neural network that computes a weighting function used for interpolation. However, in comparison to our approach, probe placement is still done manually, requiring scene-specific heuristics and/or artist input to define the precomputation volume and its spatially varying density.

Online Cache Adaptation. Techniques that precompute parts of the rendering equation struggle with fully dynamic effects, as changes in scene parameters require a new precomputation pass to capture the updated global illumination. To address this, methods have been developed to update cached quantities at runtime [MGNM19, MRNK21], transforming conventional precomputation concepts into online ones. However, to support frequent runtime updates, these methods rely on smaller caches to maintain efficiency, which compromises their representational strength. *Neural Radiance Caching* [MRNK21] mitigates this by delaying cache queries by tracing additional rays—which reduces bias and increases sampling variance—but requires additional denoising. While online cache adaptation offers robustness against arbitrary

scene changes, our work focuses on a setup that does not permit the additional cost of tracing rays or denoising.

3. Preliminaries

By excluding view-dependent lighting effects from the rendering equation [Kaj86] and focusing solely on diffuse illumination, the bidirectional reflectance distribution function (BRDF) can be factored out of the reflected radiance integral. This simplification reduces the formulation to the *irradiance* function E :

$$E(x, n) = \int_{\Omega} L_i(x, \omega_i) \langle \omega_i, n \rangle^+ d\omega_i. \quad (1)$$

At a position x , the term n can represent any direction within the spherical domain, where $E(x, \cdot)$ defines a continuous spherical function. Note that irradiance is defined for non-surface positions as well. With the diffuse BRDF being a constant proportional to the surface albedo $\rho(x)$, diffuse reflected radiance $L_r(x)$ can be simplified to:

$$L_r(x) = \frac{\rho(x)}{\pi} \cdot E(x, n). \quad (2)$$

For simplicity, we denote the *indirect* irradiance as E , with direct illumination D added at runtime. Figure 3 illustrates these quantities.

4. Neural Irradiance Volume

Our goal is to render dynamic objects that are unseen during training within known static environments using an *efficient* and *compact* representation that can be precomputed *without the need for per-scene heuristics*. We achieve this by learning a neural indirect irradiance function E_{θ} , which we term a Neural Irradiance Volume (NIV). Given a scene, we train E_{θ} via loss-driven optimization by regressing on path traced indirect irradiance values $E(x, n)$ for a dense sampling of position-direction pairs (x, n) within the scene domain. At runtime, NIV enables the introduction of unseen objects, inferring high quality diffuse illumination for both dynamic objects and static surfaces, running at around one millisecond per frame on consumer hardware on a full HD frame.

4.1. Compact Irradiance Representation

Our approach addresses the inherent inefficiencies of traditional irradiance volumes [GSHG98, MGNM19], which are typically dense and precomputed at a fixed grid resolution. This can be wasteful as a significant portion of representational capacity may be allocated inside scene geometry or in areas where the irradiance signal varies minimally. To overcome these limitations, we leverage two components: a small neural network and a multi-level hash encoding [MESK22].

Neural Network. Our model uses a four-layer fully connected coordinate-based neural network [MRNK21], with ReLU activations applied at all but the final layer, mapping an input position and direction to an irradiance value. We train a variety of neural networks, with varying overall memory capacity (refer to Table 1). The smaller models add frequency encoding [MST*21] to the input position x , mitigating the struggle of neural networks to regress on high-frequency signals [TSM*20]. The larger models make use of a learned input encoding discussed later in this section.

A notable difference compared to classic methods is how we structure our irradiance cache. Classic probes-based methods map each position x to a spherical harmonics function, which is then evaluated at a given direction ω to return irradiance E , compactly described by $x \mapsto (\omega \mapsto E)$. Our model instead maps the 5D input directly to the output, i.e. $(x, \omega) \mapsto E$. This simpler reformulation enables powerful extensions described in Section 4.2 and Section 6, and does not require the choice of a set of basis functions, like spherical harmonics.

Learned Encodings. To capture the irradiance field of larger scenes without requiring a more expensive neural network, we replace frequency encodings of our larger capacity models with a learned multi-level hash-grid encoding [MESK22]. This encoding maps positions in 3D space to latent vectors at multiple resolutions, where the latents of the finer levels are stored in a hash table. This allows adaptive use of the space, but also causes collisions of the input, as many positions may be mapped to the same latent vectors in the hash table. As demonstrated in prior work [MESK22], the collisions are implicitly handled during optimization as gradients of more important samples dominate the averaged gradient, enabling compressed representations. This prevents unimportant areas from consuming representational capacity, unlike grid-based methods where samples inside geometry would still occupy memory. On the other hand, too many collisions results in a reduction of reconstruction quality. As such, we control the size of the hash table to achieve a collision-rate that is optimal for our specific use case of representing the irradiance signal, encouraging both compactness and quality of reconstruction. Our evaluation indicates that a high collision rate only mildly affects irradiance reconstruction while severely reducing the required memory capacity, see Figure 4. NIV is parameterized with hash table size $T = 2^{17}$ for all our experiments, which strikes a balance between memory usage and representational quality on all tested scenes. The other parameters of the multi-level hash encoding are the following: the latent dimension size is 4, the scaling factor between levels is $\sqrt{2}$, the size of a side on the coarsest level is 16, and the number of levels varies between 2 and 8, depending on the target capacity (see Table 1),

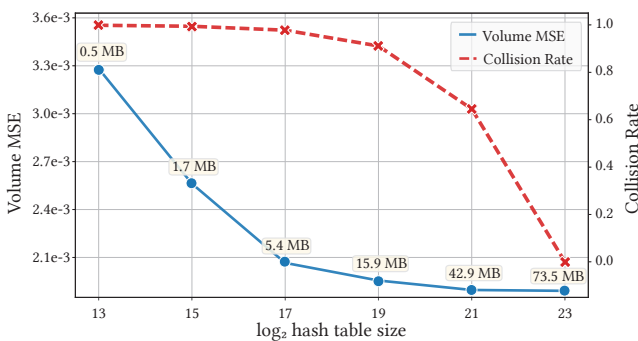


Figure 4: Varying the hash table size on Sponza (8 levels of encoding) impacts the reconstruction error of NIV. Allowing hash collisions mildly impacts the MSE while significantly reducing the required memory.

4.2. Unified Volume and Surface Cache

Irradiance probes are a practical solution because they offer a unified representation to render both dynamic and static surfaces [BB*17]. However, in practice, an additional representation is often needed because probes alone are prone to missing contact details and do not provide sufficiently high quality on static surfaces. Often this additional representation is in the form of a 2D cache dedicated to static surfaces like *light maps* [A*86], which can be precomputed to provide higher quality but comes at the cost of high memory utilization, execution divergence due to branching between light maps and irradiance probes code paths, and uv-parametrization issues.

NIV provides the advantages of both approaches—high quality on surfaces and support for unseen dynamic objects—as it inherently captures a high-quality 2D representation for static surfaces embedded within the broader 5D field. We achieve this without modifying the model architecture, but by simply adjusting the training procedure (Section 4.5): we sample a fraction of the positions explicitly on the scene’s surfaces, with their directions deterministically aligned to the corresponding surface normal. This surface domain forms a small 2D manifold of the whole 5D domain, and according to our experiments requires negligible capacity of the complete model. We experimentally verified that allocating 20% of the training samples to static surfaces strikes a good balance between volume and surface irradiance reconstruction. This targeted adjustment significantly improves the quality of rendering, effectively capturing high-frequency details such as contact shadows, without compromising the quality of dynamic objects, as shown in the false-color visualization through the banners in Sponza in Figure 6. Notably, this low-dimensional specialized representation would be infeasible to model using a grid-based cache, as grids struggle to adapt resolution to arbitrary surface manifolds without significant memory overhead or resorting to separate 2D structures with their own drawbacks.

4.3. Learning Pre-integrated Radiance

Neural rendering techniques learn a radiative quantity on the surfaces of a scene, which can be either outgoing [MRNK21, DPD22, HCZ21, CDD*24] or incoming [DKHD25] radiance. In contrast to this, our method learns a volumetric radiative quantity, the irradiance E , for each position and normal pair in the scene volume. Since we intentionally tailor our method to shade diffuse materials, we factor out the BRDF and only learn integrated incoming radiance. Using and learning a pre-integrated and directionally smooth signal enables two important properties of our method: we do not require runtime sampling and simplify the learning task.

Sampling Variance. Rendering with a neural representation of learned incident radiance $L_{i\theta}$ (i.e. not pre-integrated) is computationally expensive. Namely, at runtime $\int L_{i\theta}(x, \omega_i) \langle \omega_i, n \rangle^+ d\omega_i$ needs to be numerically evaluated which requires multiple samples such that the variance does not dominate the overall error. For high quality results, either multiple network evaluations per shading point would be needed or fewer but with additional denoising.

Representational Simplicity. Learning irradiance instead of in-

cident radiance is a more practical choice, since it is inherently smoother and less complex to represent. We validate this claim by replacing the quantity NIV learns with L_i , while maintaining the same training budget. We experimentally verify that learning such a representation takes longer to converge and significantly increases the runtime cost in the supplemental material.

Table 1: Inference times for a 4-layer network [MRNK21] driving Neural Irradiance Volume on full HD (1920×1080) frames using an RTX 4090 and i9-13900K. Networks without grid encoding (-) use eight-band positional encoding [MST*21]. “Half” renders a quarter resolution; “full” is full HD. All parameters in half-precision.

width	grid levels	full (ms)	half (ms)	memory (MB)
16	-	0.19	0.029	0.003
32	-	0.20	0.031	0.01
64	-	0.25	0.069	0.03
64	2	0.31	0.088	0.16
64	4	0.67	0.18	1.20
64	6	1.06	0.26	3.30
64	8	1.35	0.37	5.40

4.4. Rendering

After training, NIV can render frames in real-time, following a pipeline very similar to classic irradiance probes. Starting from the G-buffer of the rendered scene, the position and normal buffers are input to the neural model, which infers pixel-wise $E(x, n)$ in a single batch. Equation 2 is used to obtain indirect diffuse radiance, adding the surface emission and direct illumination for the final frame, as shown visually in Figure 1.

Half-resolution Rendering. Our method can be run at reduced resolution for higher performance, while still using the full size albedo buffer ρ for accurate materials. Running at half-resolution significantly reduces the computation time for a frame, requiring only 0.37 ms for our largest model, refer to Table 1 for the half-resolution runtime overview. Rendering at half-resolution is a common approach in real-time rendering, but introduces aliasing artifacts which could be partly resolved by using industry practice upsampling methods [KCLU07]. Since indirect irradiance varies smoothly in screen-space, aliasing is only slightly perceivable, especially after multiplication with the full-resolution albedo buffer (see supplemental video).

Dynamic Ambient Occlusion. A precomputed irradiance field, cached as either probes or NIV, does not capture how dynamic objects affect indirect illumination, but this is often not relevant in practice as added geometry is expected to be small in size compared to the scene. Some important effects might still be missed, like self-occlusion on the dynamic objects and indirect occlusion cast by them onto the scene. A common solution to mitigating part of these missed interactions is applying an ambient occlusion pass [SA07]. Since NIV already captures local high-frequency effects on static surfaces as described in Section 4.2, we run a lighter *dynamic* ambient occlusion pass, where only dynamic geometry is considered when computing the hemisphere obstruction at the shading point.

This simulates self-occlusions when shading the dynamic object, and on static surfaces it simulates how the presence of new geometry partially occludes indirect irradiance.




Figure 5: Culling training data inside scene geometry and allocating a portion of the training budget to sampling scene surfaces both improve reconstruction quality. Combining these strategies produces the most robust results.

4.5. Training

NIV is trained by uniformly sampling position-direction pairs (x, n) throughout the scene volume—20% of which are uniformly sampled on the surface geometry with their direction n set to the surface normal—and computing the associated ground-truth indirect irradiance $E(x, n)$ via path tracing. Note that while any emitter type is supported during training, real-time rendering requires light sources that can be efficiently sampled to avoid sampling noise. The model’s current output E_θ is then used to compute the relative L2 loss, normalizing the MSE by the squared network prediction [Leh18]:

$$\mathcal{L}_\theta(E(x, n), E_\theta(x, n)) = \frac{(E_\theta(x, n) - E(x, n))^2}{sg(E_\theta(x, n)^2) + \epsilon}, \quad (3)$$

where sg denotes a stop gradient operation and the constant $\epsilon = 0.01$. During training, we discard samples that are inside of surfaces, which we identify simply by checking whether the majority of normals at the first hit are backfacing whilst estimating irradiance. This avoids wasting capacity on inputs that are not visible at runtime, but also prevents mild dark leaks near static surfaces. We show the impact of sampling 20% of the training data on surfaces and culling samples with backfacing normals in Figure 5.

We use PyTorch [PGC*17] as training framework and Mitsuba3 [JSR*22] as renderer for ground truth data and Adam [Kin14] to optimize the model parameters, using a learning rate of 10^{-2} and a batch size of 2^{16} . We reduce the learning rate after the first 10k iterations by exponentially decaying to 10^{-4} , similar to other neural rendering works [HCZ21]. All our tested scenes converge after at most 50k iterations. On a single RTX 4090 a simple scene like the *Cornell Box* converges after around five minutes

or thirty minutes for a medium sized scene like *Sponza*. It is worth noting that the majority of this time is spent on path tracing irradiance (e.g. 94% of the compute budget on *Cornell Box*), as opposed to optimizing the model parameters.

5. Results

As our method is related to probe-based global illumination methods, neural surface caches, and neural rendering of variable scenes, we compare against the three families of approaches in detail.

5.1. Probe-based Methods

We compare NIV against a modern probe grid, similar to DDGI [MGNM19], implemented in our rendering system to ensure compatibility with ground truth path tracing. See supplemental material for implementation details. In our experiments, we follow the DDGI implementation but compute visibility using ray tracing rather than depth textures [MGNM19, MMNL17]. This choice favors the probe-based DDGI baseline in our quality-focused comparison, as ray-traced visibility eliminates depth-texture aliasing and avoids the additional memory overhead—where DDGI otherwise requires 1168 bytes per probe, 1024 of which are dedicated to depth buffers.

Our probe-based baseline uses second-order spherical harmonics (9 coefficients) [RH01], which are stored at half-precision using 54 bytes per probe. While further quantizing the probes to 28 bytes per probe is possible [RSS*24], a factor of 2 has little influence on the presented results even if the compression was lossless. For completeness, we also compare against the open source industry implementation of DDGI [NVI25], as shown in the supplemental material.

Memory-Error Trade-off. We evaluate the MSE on randomly sampled point-direction pairs in the scene volume, as it is a good measure for the quality for shading unseen moving objects. In all tested scenes, NIV increases quality at a given memory budget by an order of magnitude compared to probe-based methods, particularly at lower representational capacities. For a quantitative and qualitative example on *Sponza* see Figure 6. While ray-traced visibility reduces the MSE of probe-based techniques, it is still unable to close the gap to NIV and incurs a significant runtime cost.

The inferior memory scaling of probe grids is unsurprising, as the irradiance volume scales cubically with spatial discretisation and lacks the ability to adaptively allocate capacity where it is needed most. In contrast, NIV has no such constraints and can implicitly allocate capacity based on the loss function. This limitation of irradiance volumes manifests in two noticeable ways: the failure to capture contact shadows entirely and the occurrence of light leaks when ray-traced visibility is not employed, see Figure 6. Both contact shadow accuracy and leak prevention are critical for real-world applications, which the neural representation handles automatically. Additional quantitative data on other scenes can be found in the supplemental material.

Performance. The efficiency of querying a regular grid of probes is highly competitive since it consists of simple trilinear interpolation of a few spherical harmonics coefficients. The actual runtime

cost depends on the grid size and the way coefficients are laid out in memory, which influences memory transfer and caching behavior. Since our model uses a hashing function to map between grid coordinates and latent vectors, simple probe-based methods should always hold a performance edge over NIV, if both are optimized accordingly. On top, NIV requires to evaluate the input encoding and perform matrix multiplications for network evaluation [MRNK21], which altogether incurs a cost of around 1 ms on an RTX 4090. As shown in Table 1, NIV requires a minimum of 0.19 ms at full HD resolution or 0.029 ms when rendering at half resolution. Relying on a multi-resolution hash grid increases runtime cost, indicating that memory transfer costs also dominate the performance of our approach. In any case, relying on 2 or 4 grid levels (0.16 to 1.20 MB) leads to around 0.5 ms of runtime, which is easily fast enough for real-time applications and provides better quality compared to the probe-based approaches.

5.2. Neural Surface Cache

Methods that rely on a neural network to store outgoing radiance along scene surfaces [RWG*13, HCZ21, CDD*24] cannot directly shade objects that are not part of the training domain. When a previously unseen dynamic object is introduced into the scene, ray tracing is required to defer shading computations to the trained surface locations. While this approach avoids introducing additional bias, it increases variance in the final shading proportional to the deferred path depth. Caches that adapt to the scene in an online fashion [MRNK21] offer an alternative by continuously updating their neural model to account for newly introduced objects. However, both updating the cache and deferring to mitigate the bias of their (comparatively) small cache requires ray tracing during runtime.

We compare against neural surface methods by training a model similar to *Neural Radiance Caching* [MRNK21], augmented with a multi-resolution hash encoding [MESK22] with the same representational capacity as our method. After training, this model captures the outgoing radiance at any point of the static scene, but not on dynamics objects. By deferring cache queries by one bounce (or more when the deferred bounce lands on a dynamic object again) the neural surface cache can be used to estimate indirect illumination on dynamic objects. To remove the resulting variance, a denoiser must be used. Both ray tracing and denoising add a significant runtime overhead of around 5-10 ms per frame, even when using hardware supported ray tracing and highly efficient neural denoising, see Figure 8.

Results. We render scenes from our dataset with NIV and with a neural surface cache, and evaluate the rendering error with respect to path traced references. All test scenes contain dynamic objects which cover between 5% and 10% of the output image, Figure 7 shows a crop of one such dynamic object. As shown in Table 2, NIV achieves higher quality results for all inputs, while using the same memory budget and running 5-10 ms faster since ray tracing and denoising are not required. Figure 7 exemplifies the bias-variance tradeoff on the *White room* scene: the majority of artifacts from neural surface caching come from high Monte Carlo variance, hence the denoiser (Optix) does not have enough information to reconstruct an artifact-free frame. Figure 9 shows that increasing the

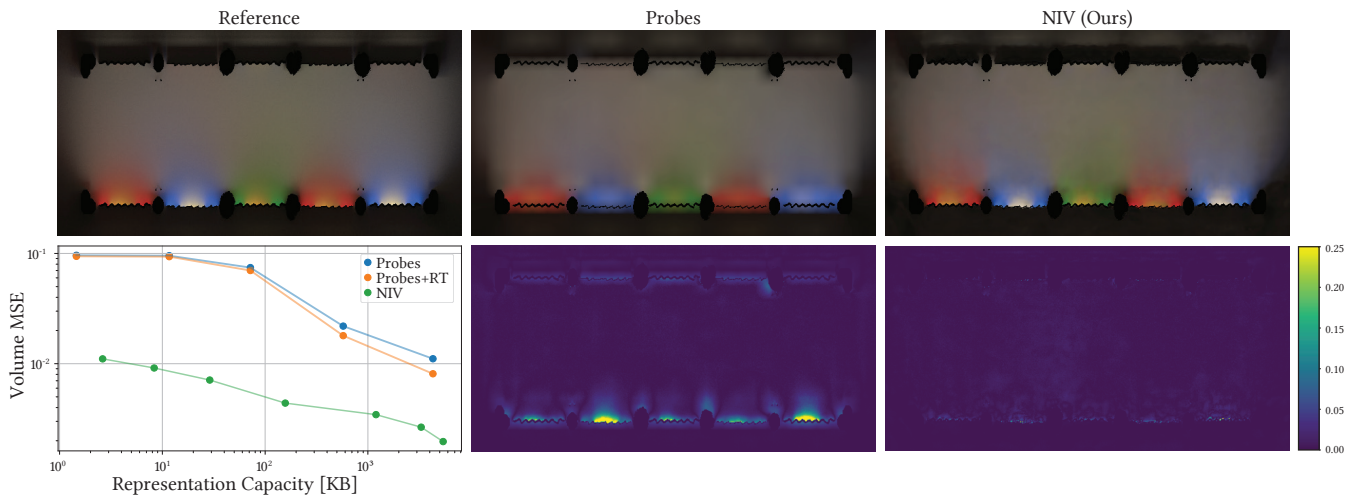


Figure 6: Horizontal slice through the banners of Sponza (5.4 MB budget). NIV captures irradiance bleed and shadows better than probe-based methods. NIV has $\sim 10\times$ higher quality across the scene’s volume. Ray tracing the visibility towards probes during evaluation (“+RT”) reduces the overall error but adds a significant performance overhead. The error map highlights the per-pixel absolute luminance error.




Figure 7: Rendering with a neural surface cache requires a single deferred surface lookup, which introduces noise (a). A neural denoiser like Optix [PBD*10] reduces noise but produces blotchiness (b, c), resolvable only with more costly ray-traced samples. Our method avoids path tracing and denoising while significantly improving quality (d) compared to a reference image (e).

surface model capacity has little effect on quality. Since variance is the major source of error, increasing the sample count is a more effective way of improving quality, but linearly raises the runtime cost.

5.3. Neural Rendering of Variable Scenes

While our goal is to shade moving objects without retraining, other methods explicitly take scene variation into account during the training phase. These approaches generally assume a predefined range for each variable scene element, which is encoded into a structured representation—usually a vector. The representation can then be used directly as an input for the neural model [DPD22], or processed through a learned encoding [SDR*24]. Such strategies attempt to learn a surface radiance field for all possible permutations of scene variables, in contrast to techniques that instead constrain learning to a single manifold within these parameters [CDD*24].

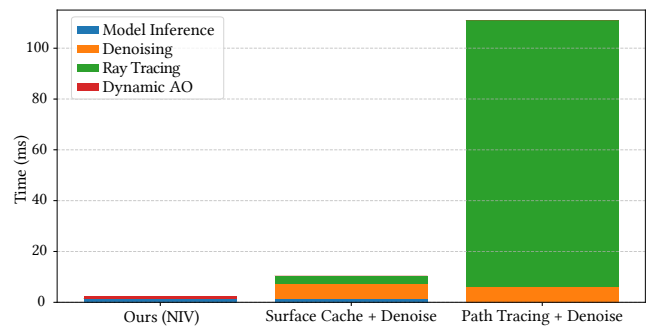


Figure 8: Runtime comparison between NIV, a denoised neural surface cache and denoised path tracing. The ray sampling budget of the methods was chosen to match the rendering quality of NIV. The surface cache only traces a single bounce (4 spp) as opposed to path tracing (16 spp), which traces full-length paths with Russian Roulette. Primary visibility is provided to all methods using a rasterization pass and is not included in the timings.

Despite correctly modeling the light transport induced by complex variable-scene interactions, these approaches present several drawbacks when compared to (neural) irradiance volumes. First, they require retraining whenever the set of scene parameters changes, such as when objects are added, removed, or modified. Second, they fail to achieve real-time performance, with inference times exceeding 100 ms in our experiments, see the supplemental material. Third, the inference time scales with the number of dynamic objects present in the scene, which limits the use of primitives with many scene variables like animated meshes. These limitations highlight challenges in deploying such methods for practical applications. Nevertheless, a natural question is how the rendering quality of these methods compares to NIV when they are similarly configured to only learn indirect diffuse global illumination.

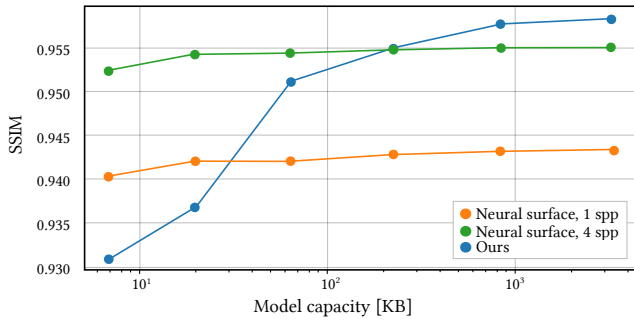



Figure 9: Rendering error of a neural surface cache and ours compared against reference on Sponza. Since the majority of the error in the neural surface caching is the variance from deferring, increasing the model capacity—and thus reducing its bias—does not improve the overall quality. Since NIV does not require sampling, its quality scales with the allocated memory.

Table 2: Comparison between a neural surface cache and NIV. Scenes from our dataset were rendered with the addition of dynamic objects and compared to path traced references. Both methods used a capacity of 5.40 MB. We use the HDR variant of FLIP [ANSA21].

Scene	FLIP (\downarrow)		MSE (\downarrow)	
	Surface	Ours	Surface	Ours
Sponza	0.25	0.10	2.32e-3	9.09e-6
Cornell box	0.14	0.05	5.57e-4	2.53e-5
Dining room	0.27	0.24	3.07e-2	3.31e-3
Bathroom	0.15	0.06	3.22e+0	2.76e-3
White room	0.17	0.09	2.76e-3	8.13e-4
Living room	0.10	0.04	1.14e-2	2.36e-4

NIV does not explicitly model higher-order interactions between dynamic objects and the surrounding scene. In contrast, methods based on variable scene encodings need to devote part of their representational capacity to learning these interactions on top of learning the outgoing radiance on the static scene geometry. In our results we find that for a fixed model size, NIV often achieves comparable or lower rendering error, particularly as the number of scene variables increases.

As illustrated in Figure 10 and Figure 11, the rendering error achieved by NIV lies within the same range as methods that were trained with explicit knowledge of scene objects. This suggests that the irradiance volume formulation, combined with dynamic ambient occlusion, is highly effective at generalizing to unseen configurations. However, there exists a theoretical lower bound to irradiance volumes in this setup, represented in Figure 11 by the dotted line. This bound corresponds to the error level of a path-traced static irradiance field computed without the dynamic objects. As the model capacity of variable-scene methods grows, their rendering error can drop below this threshold, allowing them to produce shading effects that more closely resemble path traced global illumination, including the higher-order interactions between the variable parameters and the scene. In the supplemental material, we show an additional figure similar to Figure 11, where the variable



(a) 5 Variable Parameters

(b) 35 Variable Parameters

Figure 10: At a model capacity of 1 MB NIV captures color bleed from the static scene onto the variable scene elements—the floating armadillos—while the variable scene methods [DPD22, SDR*24] do not. As opposed to the compared methods, NIV has not seen dynamic objects during training, and can be queried with any amount of them without re-training. See Figure 11 for quantitative data.

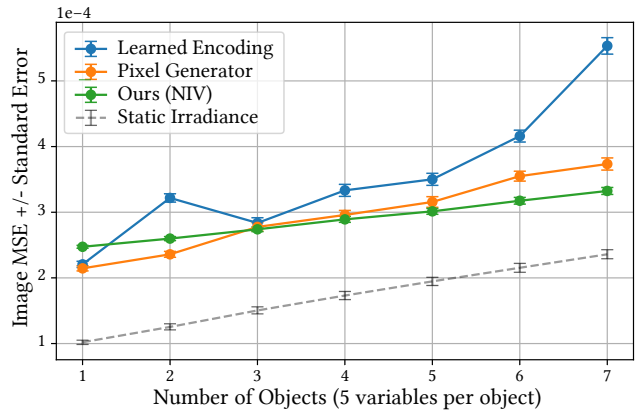


Figure 11: We compare NIV with variable-scene methods that encode changing scene parameters via a PixelGenerator [DPD22] or a learned encoding [SDR*24]. Unlike these approaches, which require training for each scene configuration, NIV is trained only once on the static irradiance of the empty scene. Yet it achieves lower rendering error at higher object counts, underlining the effectiveness of the irradiance assumption. All methods use the same 1 MB capacity, with each object adding five degrees of freedom (35 scene variables at the rightmost data point).

scene methods are ran using the recommended model capacities as described in the respective works [DPD22, SDR*24], which improves their rendering quality at the cost of extra model capacity and an even slower runtime.

6. Application: Higher-dimensional Irradiance Fields

In many real-world applications, only a known set of scene parameters is expected to change at runtime [SSS*20], e.g. rotating emitters that mimic a time-of-day cycle. Given the prevalence of such behavior we extend pre-computing a neural irradiance volume with respect to such changing scene parameters, by introducing new in-

puts to the model, e.g. the current position of the moving emitters. We train the irradiance field while randomly sampling the variable scene parameter, similar to the neural variable scene methods in Section 5.3. After training, this removes the need to update the representation in real-time and avoids potential temporal artifacts and updating costs.

Training with many variable parameters, as highlighted before, significantly impacts the performance of a neural representation. To maintain performance we targeted learning 1-2 variables that model dynamic phenomena useful in practice. We find it to be sufficient to apply frequency encoding [MST*21] to such changes, which only adds little inference overhead. Scaling to more than two scene variables at the same representational capacity causes visible reconstruction error, which could be resolved through the use of learned encodings, with the same disadvantages shown in Section 5.3. We implemented two use cases: a rotating directional light source simulating a time-of-day cycle and a moving occluding object in the scene. See results in Figure 12 and the supplemental video. To ensure artifact-free results, we double the size of the hash table that encodes the scene position for this application (i.e. $\sim 10\text{MB}$).




Figure 12: Adding the angle of a directional light to NIV lets us model time-of-day changes in *Sponza*—a common use case in game development—without requiring further training at runtime. The Bunny is unseen during training, but receives high-quality color bleed from the scene.

7. Limitations & Future Work

Direct Illumination. A fundamental assumption of NIV is that direct illumination can be estimated efficiently at runtime without introducing sampling-related noise. However, this assumption only holds for simple lighting and low numbers of lights, and may require additional render passes, e.g., shadow maps [Wil78]. While many real-time techniques exist to handle the many lights problem [BWP*20], they often exhibit significant noise levels themselves and thus efficient noise-free shading is impossible.

Irradiance Volume Assumptions. Traditional probe based illumination assumes the impact of moving objects on the scene’s light transport is small—or artist controlled [SSS*20]—which holds true for many real-life applications and does not hinder its widespread use in games [BB*17]. We inherit this assumption, and do not capture the moving object’s impact on the scene’s light transport. As with probe based illumination, extensions would be needed for glossy materials [RLP*20] or higher-order self-occlusions [SKS02].

Compression. Further compressing multi-level hash encoding

and MLP architectures is an active research topic [TMND*23]. Exploiting such compression might further improve the memory-error trade-off and runtime cost of NIV. It is worth noting that such compression can introduce artifacts that might require problem-specific regularization terms or training parameter tuning [DMD*23]. Our early experiments show that using a strong Laplacian regularizer and a lower learning rate significantly helps reconstruction.

Glossy materials. NIV can be used with non-diffuse materials by deferring queries to diffuse intersections, at the cost of sampling-related noise. To maintain noise-free rendering, alternative solutions could be explored. A promising direction incorporates surface roughness as a control parameter in the neural model [VHM*22], possibly allowing the method to natively handle glossy materials.

Online Learning. Updating a learned cache at runtime has been shown to be possible for surface-based representations [MRNK21]. Extending NIV with online learning works, but is inherently slower to learn since it has a larger input domain—the scene volume, as opposed to a subset of the scene surfaces—and requires a low bias representation as it avoids deferring queries to other scene locations. We experimentally verified that starting from a learned NIV significantly improves convergence as opposed to learning from scratch. However, to update it in real-time, further extensions are required such as only sampling the camera’s view frustum and/or employing loss-driven sampling [DPD22].

Larger Scenes. When scaling NIV to larger scenes, we found that model capacity eventually becomes a limiting factor, leading to increased reconstruction error. One way to address this is to partition the scene into a grid of smaller regions [RPLG21] and learn a different NIV for each of them, which allows the method to retain accuracy by distributing representational capacity. An interesting direction for future work is striking a balance between traditional spatial subdivision and the neural network capacity as explored in related work [WRM*24]. Another promising avenue is the integration of level-of-detail techniques, allocating capacity adaptively or selecting a lower frequency model based on the position of the camera.

8. Conclusion

We introduced a Neural Irradiance Volume (NIV), a novel approach that modernizes real-time global illumination for dynamic objects. By compressing the indirect irradiance field in a compact neural model, NIV overcomes the limitations of traditional probe-based methods, offering superior visual quality and an order of magnitude improvement in memory efficiency. Our technique provides a unified, noise-free, and high-quality solution that is practical for real-time applications. By eliminating the need for expensive runtime operations such as ray tracing or denoising, which are common in neural rendering techniques, NIV shows that high-quality indirect lighting is achievable without sacrificing performance. We believe that NIV has the potential to benefit existing applications that utilize irradiance volumes, such as game engines, by reducing memory requirements and complexity, and improving visual quality. Furthermore, its differentiable nature makes it a promising building block for inverse rendering pipelines and other neural scene representations.




Figure 13: Overview of the indirect irradiance—i.e. without albedo for clarity—using our Neural Irradiance Volume. The movable objects *Lucy*, *Bunny* and *Dragon* are placed in the scene at render time and are not seen during training. Since NIV and probe volumes capture only the indirect irradiance of the known scene, dynamic ambient occlusion is also added as described in Section 4.4. All insets can be compared with a high sample count path traced reference. Compared to NIV, probe grids noticeably struggle with several kinds of light leaks and incorrect color bleed.

References

- [A*86] ARVO J., ET AL.: Backward ray tracing. In *Developments in Ray Tracing, Computer Graphics, Proc. of ACM SIGGRAPH 86 Course Notes* (1986), pp. 259–263. 4
- [ANSA21] ANDERSSON P., NILSSON J., SHIRLEY P., AKENINE-MÖLLER T.: Visualizing Errors in Rendered High Dynamic Range Images. In *Eurographics Short Papers* (May 2021). 8
- [BB*17] BARRÉ-BRISEBOIS C., ET AL.: A certain slant of light: Past, present and future challenges of global illumination in games. *ACM SIGGRAPH Courses* (2017). 2, 4, 9
- [BWP*20] BITTERLI B., WYMAN C., PHARR M., SHIRLEY P., LEFOHN A., JAROSZ W.: Spatiotemporal reservoir resampling for real-time ray tracing with dynamic direct lighting. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 148–1. 9
- [CDD*24] COOMANS A., DOMINICI E. A., DÖRING C., MUELLER J. H., HLADKY J., STEINBERGER M.: Real-time neural rendering of dynamic light fields. *Computer Graphics Forum* 43, 2 (2024), e15014. 2, 3, 4, 6, 7
- [DKHD25] DEREVIANNYKH M., KLEPIKOV D., HANIKA J., DACHS-BACHER C.: Neural two-level monte carlo real-time rendering. In *Computer Graphics Forum* (2025), Wiley Online Library, p. e70050. 4
- [DMD*23] DATTA S., MARSHALL C., DONG Z., LI Z., NOWROUZEZHRAI D.: Efficient graphics representation with differentiable indirection. In *SIGGRAPH Asia 2023 Conference Papers* (2023), pp. 1–10. 9
- [DPD22] DIOLATZIS S., PHILIP J., DRETTAKIS G.: Active exploration for neural global illumination of variable scenes. *ACM Transactions on Graphics (TOG)* 41, 5 (2022), 1–18. 2, 3, 4, 7, 8, 9
- [GSHG98] GREGER G., SHIRLEY P., HUBBARD P. M., GREENBERG D. P.: The irradiance volume. *IEEE Computer Graphics and Applications* 18, 2 (1998), 32–43. 2, 3
- [HCZ21] HADADAN S., CHEN S., ZWICKER M.: Neural radiosity. *ACM Transactions on Graphics (TOG)* 40, 6 (2021), 1–11. 2, 3, 4, 5, 6
- [Hoo16] HOOKER J. T.: Volumetric global illumination at treyarch. *ACM SIGGRAPH Advances in Real-Time Rendering Course*. (2016). 2
- [IS17] IWANICKI M., SLOAN P.-P.: Precomputed lighting in call of duty: Infinite warfare. *Advances in Real-Time Rendering, Part I (ACM SIGGRAPH 2017 Courses)*, Association for Computing Machinery, New York, NY, Article 7a (2017). 2
- [ISS24] IWANICKI M., SLOAN P.-P., SILVENNOINEN A., SHIRLEY P.: Neural light grid: Modernizing irradiance volumes with machine learning. *ACM SIGGRAPH Advances in Real-Time Rendering in Games Course* (2024). 3
- [JSR*22] JAKOB W., SPEIERER S., ROUSSEL N., NIMIER-DAVID M., VICINI D., ZELTNER T., NICOLET B., CRESPO M., LEROY V., ZHANG Z.: Mitsuba 3 renderer, 2022. URL: <https://mitsuba-renderer.org>. 5
- [Kaj86] KAJIYA J. T.: The rendering equation. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques* (1986), pp. 143–150. 2, 3
- [KCLU07] KOPF J., COHEN M. F., LISCHINSKI D., UYTENDAELE M.: Joint bilateral upsampling. *ACM Transactions on Graphics (TOG)* 26, 3 (2007), 96–es. 5
- [Kin14] KINGMA D. P.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014). 5
- [Leh18] LEHTINEN J.: Noise2noise: Learning image restoration without clean data. *arXiv preprint arXiv:1803.04189* (2018). 5
- [MESK22] MÜLLER T., EVANS A., SCHIED C., KELLER A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)* 41, 4 (2022), 1–15. 3, 4, 6
- [MGNM19] MAJERCIK Z., GUERTIN J.-P., NOWROUZEZHRAI D., MCGUIRE M.: Dynamic diffuse global illumination with ray-traced irradiance fields. *Journal of Computer Graphics Techniques* 8, 2 (2019). 3, 6
- [MMNL17] MCGUIRE M., MARA M., NOWROUZEZHRAI D., LUEBKE D.: Real-time global illumination using precomputed light field probes. In *Proceedings of the 21st ACM SIGGRAPH symposium on interactive 3D graphics and games* (2017), pp. 1–11. 2, 6
- [MRNK21] MÜLLER T., ROUSSELLE F., NOVÁK J., KELLER A.: Real-time neural radiance caching for path tracing. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–16. 2, 3, 4, 5, 6, 9
- [MST*21] MILDENHALL B., SRINIVASAN P. P., TANCIC M., BARRON J. T., RAMAMOORTHY R., NG R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* 65, 1 (2021), 99–106. 3, 5, 9
- [NVI25] NVIDIA GAMEWORKS: Rtxgi-ddgi. <https://github.com/NVIDIAGameWorks/RTXGI-DDGI>, 2025. 6
- [PBD*10] PARKER S. G., BIGLER J., DIETRICH A., FRIEDRICH H., HOBEROCK J., LUEBKE D., MCALLISTER D., MCGUIRE M., MORLEY K., ROBISON A., ET AL.: Optix: a general purpose ray tracing engine. *Acm transactions on graphics (tog)* 29, 4 (2010), 1–13. 7
- [PGC*17] PASZKE A., GROSS S., CHINTALA S., CHANAN G., YANG E., DEVITO Z., LIN Z., DESMAISON A., ANTIGA L., LERER A.: Automatic differentiation in pytorch. In *NIPS-W* (2017). 5
- [RBRD22] RAINER G., BOUSSEAU A., RITSCHEL T., DRETTAKIS G.: Neural precomputed radiance transfer. *Computer Graphics Forum* 41, 2 (2022), 365–378. 3
- [RH01] RAMAMOORTHY R., HANRAHAN P.: An efficient representation for irradiance environment maps. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (2001), pp. 497–500. 6
- [RHP*24] REN H., HUO Y., PENG Y., SHENG H., XUE W., HUANG H., LAN J., WANG R., BAO H.: Lightformer: Light-oriented global neural rendering in dynamic scene. *ACM Transactions on Graphics (TOG)* 43, 4 (2024), 1–14. 3
- [RPL*20] RODRIGUEZ S., LEIMKÜHLER T., PRAKASH S., WYMAN C., SHIRLEY P., DRETTAKIS G.: Glossy probe reprojection for interactive global illumination. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–16. 2, 9
- [RPLG21] REISER C., PENG S., LIAO Y., GEIGER A.: Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *Proceedings of the IEEE/CVF international conference on computer vision* (2021), pp. 14335–14345. 9
- [RSS*24] ROUGHTON T., SLOAN P.-P., SILVENNOINEN A., IWANICKI M., SHIRLEY P.: Zh3: Quadratic zonal harmonics. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 7, 1 (2024), 1–15. 6
- [RWG*13] REN P., WANG J., GONG M., LIN S., TONG X., GUO B.: Global illumination with radiance regression functions. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 130–1. 2, 3, 6
- [SA07] SHANMUGAM P., ARIKAN O.: Hardware accelerated ambient occlusion techniques on gpu. In *Proceedings of the 2007 symposium on Interactive 3D graphics and games* (2007), pp. 73–80. 5
- [SDR*24] SU R., DONG H., REN J., JIN H., CHEN Y., WANG G., LI S.: Dynamic neural radiosity with multi-grid decomposition. In *SIGGRAPH Asia 2024 Conference Papers* (2024), pp. 1–12. 2, 3, 7, 8
- [SKS02] SLOAN P.-P., KAUTZ J., SNYDER J.: Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics (TOG)* 21, 3 (jul 2002), 527–536. 2, 9
- [SSS*20] SEYB D., SLOAN P.-P., SILVENNOINEN A., IWANICKI M., JAROSZ W.: The design and evolution of the UberBake light baking system. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 39, 4 (July 2020). 2, 8, 9

- [ST15] SILVENNOINEN A., TIMONEN V.: Multi-scale global illumination in quantum break. *ACM SIGGRAPH Advances in Real-Time Rendering Course*. (2015). 2
- [TMND*23] TAKIKAWA T., MÜLLER T., NIMIER-DAVID M., EVANS A., FIDLER S., JACOBSON A., KELLER A.: Compact neural graphics primitives with learned hash probing. In *SIGGRAPH Asia 2023 Conference Papers* (2023), pp. 1–10. 9
- [TSM*20] TANCİK M., SRINIVASAN P., MILDENHALL B., FRIDOVICH-KEIL S., RAGHAVAN N., SINGHAL U., RAMAMOORTHI R., BARRON J., NG R.: Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in neural information processing systems* 33 (2020), 7537–7547. 3
- [VHM*22] VERBIN D., HEDMAN P., MILDENHALL B., ZICKLER T., BARRON J. T., SRINIVASAN P. P.: Ref-nerf: Structured view-dependent appearance for neural radiance fields. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022), IEEE, pp. 5481–5490. 9
- [VPG14] VARDIS K., PAPAIOANNOU G., GKARAVELIS A.: Real-time radiance caching using chrominance compression. *Journal of Computer Graphics Techniques (JCGT)* 3, 4 (December 2014), 111–131. URL: <http://jcgt.org/published/0003/04/06/>. 2
- [Wil78] WILLIAMS L.: Casting curved shadows on curved surfaces. In *Proceedings of the 5th annual conference on Computer graphics and interactive techniques* (1978), pp. 270–274. 9
- [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 Interactive 3D Graphics and Games* (New York, NY, USA, 2019), I3D '19, Association for Computing Machinery. 2
- [WRM*24] WEIER P., RATH A., MICHEL É., GEORGIEV I., SLUSALLEK P., BOUBEKEUR T.: N-bvh: Neural ray queries with bounding volume hierarchies. In *ACM SIGGRAPH 2024 Conference Papers* (2024), pp. 1–11. 9
- [ZDP*25] ZENG C., DONG Y., PEERS P., WU H., TONG X.: Renderformer: Transformer-based neural rendering of triangle meshes with global illumination. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers* (2025), pp. 1–11. 3
- [ZLZ*25] ZHOU Z., LI C., ZHANG Z., TANG M., LI Z., LUAN S., HUANG Z.: Gaussian compression for precomputed indirect illumination. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers* (2025), pp. 1–10. 2
- [ZWZ*25] ZHU J., WU Z., ZHANG Q., LIAO C., HUANG Z.: Wishgi: Lightweight static global illumination baking via spherical harmonics fitting. *ACM Transactions on Graphics (TOG)* 44, 4 (2025), 1–12. 2