

Neural Scene Baking for Permutation Invariant Transparency Rendering with Real-time Global Illumination

Ziyang Zhang(✉)¹ and Edgar Simo-Serra¹

© The Author(s)

Abstract Neural rendering provides a fundamentally new way to render photorealistic images. Similar to traditional light-baking methods, neural rendering utilizes neural networks to bake representations of scenes, materials, and lights into latent vectors learned from path-tracing ground truths. However, existing neural rendering algorithms typically use G-buffers to provide position, normal, and texture information of scenes, which are prone to occlusion by transparent surfaces, leading to distortions and loss of detail in the rendered images. To address this limitation, we propose a novel neural rendering pipeline that accurately renders the scene behind transparent surfaces with global illumination and variable scenes. Our method separates the G-buffers of opaque and transparent objects, retaining G-buffer information behind transparent objects. Additionally, to render the transparent objects with permutation invariance, we designed a new permutation-invariant neural blending function. We integrate our algorithm into an efficient custom renderer to achieve real-time performance. Our results show that our method is capable of rendering photorealistic images with variable scenes and viewpoints, accurately capturing complex transparent structures along with global illumination. Our renderer can achieve real-time performance (256×256 at 63 FPS and 512×512 at 32 FPS) on scenes with multiple variable transparent objects.

Keywords global illumination; neural rendering; real-time rendering; transparency rendering

1 Introduction

Global illumination rendering has been a fundamental challenge in computer graphics due to the complexity of the physical phenomena involved. Traditionally, path tracing

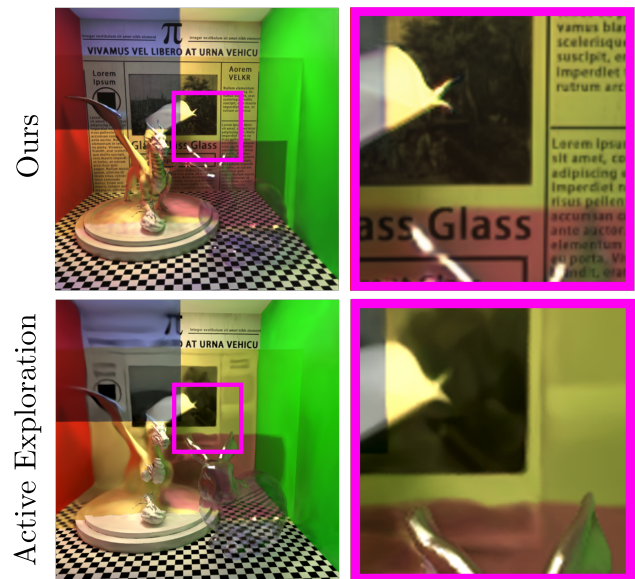


Fig. 1 Transparency quality improvement. We introduce our custom neural rendering framework to address transparency rendering issues in existing neural rendering approaches. Compared to Active Exploration, our method is capable of rendering objects behind transparent surfaces with minimum loss of details with global illumination effects.

algorithms [1, 2] have been the main method to render photorealistic images with high-quality global illumination effects. However, such methods are known to be expensive and time-consuming; rendering one image may take minutes. Baked global illumination techniques, on the other hand, provide efficient solutions for interactive rendering by pre-computing lighting information. Methods such as lightmap baking and light probes [3] have been commonly used in game engines to generate real-time global illumination, but they have limitations on dynamic scenes.

Recently, neural rendering approaches have emerged as promising alternative for realistic image synthesis. Using neural networks to represent scenes, materials, and lighting can be seen as one type of pre-computation method. In the

¹ Faculty of Science and Engineering, Waseda University, Tokyo, Japan. E-mail: Z. Zhang(✉), ziyangz5@toki.waseda.jp; E. Simo-Serra, ess@waseda.jp

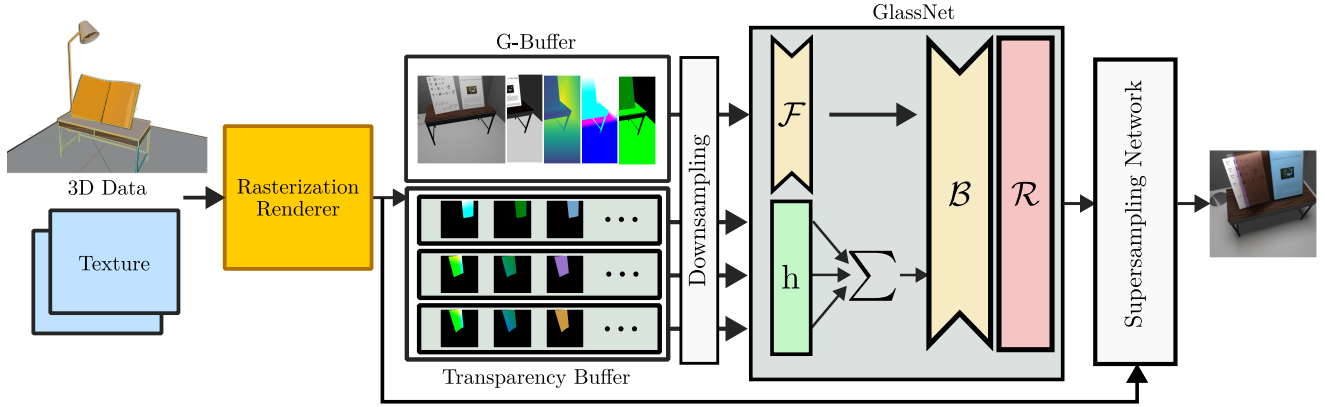


Fig. 2 Overview of our neural scene rendering framework. First, our rasterization-based renderer will first render the G-buffer, direct lighting, and transparency buffers. Then, a neural network, which we denote as *GlassNet*, will use those buffers and rendering results as inputs to synthesize high quality images with global illumination and accurate transparency. Details of *GlassNet* can be found in Sec. 3.4 and Fig. 3.

early stage, neural rendering methods were only capable of rendering fixed scenes under limited types of lighting with complex spatial inputs [4]. More recently, by using G-buffer as inputs, newer methods avoid using complex spatial data such as point clouds, and are able to render scenes with variable geometries, materials, and lighting [5–9]. However, the usages of G-buffer re-introduces the classic problem of transparency rendering. When generating G-buffers, the shading information will be overwritten by transparent surfaces in the front, as shown in Fig. 4. Without such crucial shading information, rendering models can only hard code all the information into the neural network based on ground truths, which is unrealistic and ineffective, leading to poor quality.

To address this limitation, we propose a novel neural rendering framework that separately renders opaque objects and semi-transparent surfaces into two groups of G-buffers. By doing this, all necessary shading information of both opaque and transparent surfaces are preserved and can be used for further rendering. Furthermore, we design a neural blending function to encode the transparency buffer into a latent representation with permutation invariance to achieve Order-Independent-Transparency (OIT). Finally, we use a deep CNN (Convolutional Neural Network) to encode scenes and blend them with transparency representations, rendering the resulting image with global illumination.

Our results show that we significantly improve the transparency quality, exceeding the most recent neural approaches, Active Exploration [8] and Neural Global Illumination [9]. Our method can accurately render the geometry and texture details under transparent surfaces, whereas methods without transparency buffers converge to low-quality results.

In summary, our main contributions are:

- (1) A novel neural rendering framework preserves both opaque and transparent information simultaneously on buffers, providing accurate shading information on rendering transparency.
- (2) A neural blending function that encodes transparent surfaces with permutation invariant, achieving OIT.
- (3) A real-time renderer capable of rendering images with high performance enabled by the above algorithms.

2 Related Work

2.1 Traditional Transparency Rendering

In the traditional forward rendering pipeline, various techniques have been used to handle transparency. One widely used algorithm is alpha blending [10], which utilizes a separate alpha channel to control the color composition of semi-transparent objects with opaque objects on each fragment. However, alpha blending is order-dependent and requires sorting of objects by depth, which can be impractical. The high computational complexity is not the only problem of those sorting-based algorithms. When transparent objects intersect with each other, it is impossible to sort the objects globally unless the sorting is done by fragments, which again greatly increases the computational complexity. Another algorithm called the A-Buffer algorithm [11] stores a list of relevant fragments in each pixel, but sorting of all fragments is required with unbounded memory usage. The Z^3 algorithm [12] is a similar algorithm with bounded memory usage but comes with a trade-off in image quality. In contrast, our proposed method does not require any sorting, and does not generate a large amount of frame buffers during rendering procedure.

To avoid sorting, several techniques have been introduced.

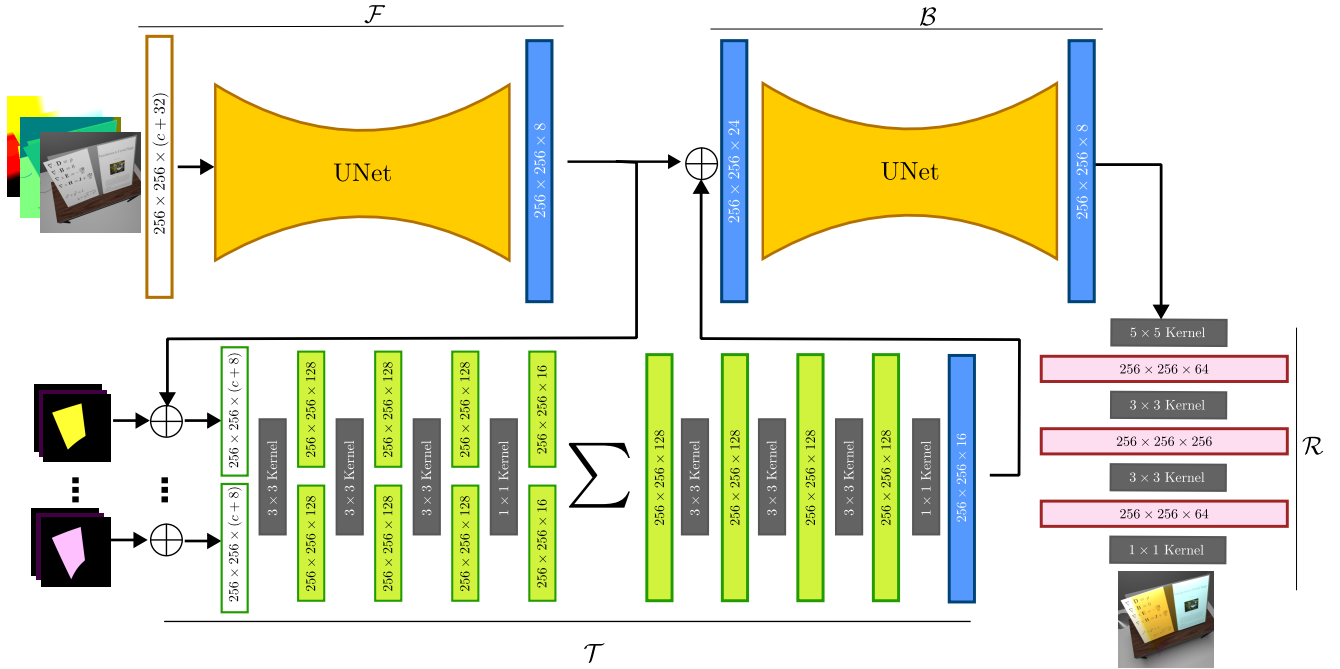


Fig. 3 Architecture of GlassNet. Our proposed method, GlassNet, contains four building blocks including the scene encoder, \mathcal{F} ; the permutation invariant transparency buffer blending function, \mathcal{T} ; the final blending network, \mathcal{B} ; and the rendering network, \mathcal{R} . All parts are trained jointly.

Depth peeling [13] renders objects in multiple passes, peeling transparent objects based on the z-buffer. However, the number of passes required by depth peeling increases the computational cost.

These algorithms often encounter challenges in the modern deferred rendering pipeline. Deferred rendering [14] is a technique that separates the geometric stage and lighting stage into two passes. In the geometry pass, all the geometry information is rendered into G-buffers to be used by the lighting pass for shading. While deferred rendering significantly improves rendering performance with multiple light sources, it introduces complexity to transparency rendering. Rendering geometry information into the G-buffer causes the objects behind transparent surfaces to be occluded, making it impossible to retain all geometry information without a prohibitively large number of buffers. Several algorithms solve this problem with worsened image quality. [15, 16]. By applying a permutation-invariant blending algorithm, our method can combine generated transparency information during inference, while only requiring two frame buffers regardless of the number of transparent objects, significantly lowering the memory usage.

Several neural rendering approaches utilize the G-buffers as input for their neural networks [6, 8, 9, 17], resulting in these methods also encountering transparency occlusion issues with the G-buffers, causing bad image quality. In contrast,

our method is not affected by such issues and supports OIT rendering.

2.2 Baked Global Illumination

Different from expensive path tracing methods, which are the main approach to render photorealistic images in the modern film industry [18], baked or precomputed global illumination methods are more often used in interactive or real-time settings [19].

Lightmap baking is a popular technique used widely in modern game engines [20]. It precomputes the diffuse global illumination and uses interpolation to generate global illumination during rendering. Several methods have been proposed to generate the lightmap offline, such as by radiosity [21] or path-tracing. Naively, baked lightmap only supports static scenes with diffuse objects. Light probes technique was introduced to enable global illumination for dynamic objects [3]. McGuire *et al.* [22] extended light probes to support glossy surfaces. However, such methods are limited to either static lighting or a limited amount of light transportation types.

Several neural approaches use neural networks to encode the scene representation into neural vectors and synthesize global illumination such as Compositional Neural Scene Representations (CNSR) [6], Active Exploration and Neural Global Illumination. Similar to such neural approaches, our method can render global illumination with dynamic scenes

with various light transportation types, with support for transparent objects.

2.3 Neural Rendering

Using neural networks for rendering is a rapidly developing research topic that provides fundamentally new approaches to photorealistic rendering. We will introduce the closest related paper to this work. We refer to [23] for a more detailed neural rendering survey.

Initially, Ren *et al.* [4] used neural networks to learn indirect lights from dynamic lighting and roughness with fixed geometry, complex spatial data, and only point lights. Later, Eslami *et al.* [5] proposed a method that trains an encoder-decoder neural network to represent scenes as latent vectors. CNSR by Granskog *et al.* further improved the interpretability of the neural scene representation by partitioning the scene representation vectors into sections of lighting, materials, geometry, etc. Similar to our method, most recent solutions [6–9] use G-buffers as the input to the neural networks to produce indirect lighting and global illumination. Using G-buffers instead of point clouds or voxel data simplifies the input data. Such algorithms can usually be seen as precomputed light transportation and a form of scene baking. Using neural networks to precompute wavelet function has also shown strong improvement on glossy global illumination [24]. Attempts on high performance real-time neural rendering are also made by Xin *et al.* [17], but only limited on single-bounce diffuse indirect illumination

However, using G-buffers as inputs brings back the classic issue caused by transparent objects. The G-buffers cannot provide information on objects behind transparent surfaces, ending with poor rendering quality by such methods. In our framework, we use independent G-buffers for the transparent objects and a permutation invariant neural network to achieve OIT. Note that the recent studies in neural rendering are mainly orthogonal to our proposed method with different goals. Our method can be directly combined with those methods to improve the rendering quality further.

2.4 Permutation Invariant Machine Learning

PointNet [25] uses a multi-layer perceptron network to transform the input into latent vectors and uses a symmetric function to compose such vectors, which is proved to be a simple but strong method. PointNet inspires how our rendering neural network can process each transparency input to cancel the permutation order dependency. We refer a more detailed review of permutation-invariant neural network to Section 6 of [26].

3 Proposed Approach

3.1 Overview

Given a scene, the goal of our approach is to bake a neural scene representation into a neural network that can correctly approximate global illumination with variable parameters of material, geometry, and lighting without being affected by transparent surfaces. For any shading position, x , on surfaces, the outgoing radiance, L_o , can be calculated by the Rendering Equation[1]:

$$L_o(x, \omega_o) = \int_{\Omega} L_i(x, \omega_i) f_x(x, \omega_o, \omega_i) |n_x \cdot \omega_i| d\omega_i + L_e(x, \omega_o) \quad (1)$$

where f_x and n_x represent the BSDF and normal; ω_o indicates outgoing or viewing direction, ω_i indicates the incident direction; L_i and L_e denote the incident radiance and emission radiance. Ω denotes a hemisphere.

The Rendering Equation can also be written into the following form such as the direct lighting and the indirect lighting are separated:

$$L_o(x, \omega_o) = L_e(x, \omega_o) + L_d(x, \omega_o) + L_g(x, \omega_i) \quad (2)$$

where L_d is the direct lighting, and L_g is the indirect lighting which contributes mainly to the global illumination. Current methods use high-dimension complex non-linear functions, *i.e.*, neural networks, to predict $L_o(x, \omega_o)$ given all the necessary geometry, lighting, and material parameters as inputs. In contrast, our approach consists of a neural network, which we denote as *GlassNet*, to predict the indirect lighting, L_g , as shown in Fig. 2. We use Linearly Transformed Cosines (LTCs) [27] as a sufficient approximation to the direct illumination term L_d as an input to our neural network to achieve efficiency. Notice that as an input to the neural network, L_d is not necessarily to be analytically accurate because the neural network will learn the complexity of the lighting transportation from the path tracing ground truths.

Most recent methods provide necessary shading information such as ω_o , position, texture, and material parameters through G-buffers. However, the naive approach using G-buffers as input brings a significant problem: as shown in Fig. 4, important shading information is hidden by transparent surfaces. We propose a new neural rendering framework that separately draws opaque G-buffers and transparency buffers, preserving all the information, including the transparent surfaces and the objects behind those. Then, we introduce a symmetric neural function to cancel the permutation dependency in the transparency buffers and generate a neural transparency representations as shown in Sec. 3.3.




Fig. 4 Our G-buffer scene representation approach. Instead of naively rendering the entire scene to a single set of G-buffers, we use separate buffers for the transparent objects, allowing us to represent complex transparency visibility in scenes accurately.

3.2 Transparency Buffers

For a given scene, we rasterize each transparent surface into a list of G-buffers containing all the information required by a traditional path tracer to render L_o , including position, surface normal, albedo, ω_o , and material parameters as shown in Fig. 5. We generate a set of buffers with a size of $(w, h, c * (t + 1))$, where w , h and c denote the width, height and channel of images. t is the number of transparent objects in the scene. Finally, our rasterization-based renderer creates the opaque G-buffers, and use LTCs to compute the direct lighting L_d . By doing so, we can preserve all the necessary information on both normal surfaces and transparent surfaces.

However, further processing is necessary for the naive transparency buffer to be a good representation as a neural network input. Generally, neural networks are known to depend on the order of inputs[25]. Inputting the naive transparency buffer will cause our rendering network to be order-dependent on transparency. With modern rendering engine optimization features such as batch rendering[28] and occlusion culling[29], the order of rendered objects cannot be guaranteed. The surfaces can even be culled during rendering. It is also infeasible to sort the rendered buffers after rendering. In order to solve this problem, we designed a symmetric neural network invariant to permutation to blend the transparency buffer.

3.3 Permutation Invariant Neural Representation

We design a blending invariant algorithm to cancel any sensitivity to the order of the inputs while significantly lowering the memory usage during inference as shown in Sec. 4.4.1, inspired by traditional transparency algorithms such as alpha blending. In alpha blending, the function, T , of compositing

transparent objects can be described as[10, 30, 31]:

$$T\left(\{(C_1, \alpha_1), (C_2, \alpha_2), \dots, (C_t, \alpha_t)\}\right) = \sum_{i=1}^t \alpha_i C_i z_i \quad (3)$$

where C_i and α_i respectively are the color and the alpha value of each transparent object. z_i is the product of all alpha values in front of the object:

$$z_i = \begin{cases} 1 & \text{if } i = t \\ \prod_{j=1}^{t-i} (1 - \alpha_j) & \text{otherwise} \end{cases} \quad (4)$$

This function is clearly input-order-dependent, but it can be revised into a function with permutation invariant to inputs.

Based on the alpha blending function with the inspiration from the PointNet, we design our permutation invariant transparency buffer blending function, \mathcal{T} , to generate the neural representation, τ :

$$\begin{aligned} \mathcal{T}(\{b_1, b_2, \dots, b_t\}, \sigma) \\ \approx g(h(b_1, \sigma), h(b_2, \sigma), \dots, h(b_t, \sigma)) \\ = \sum_{i=1}^t h(b_i, \sigma) = \tau \end{aligned} \quad (5)$$

where b is the transparency buffer, g is a symmetric function, σ is a scene representation, and h is a neural network shared by all the transparency buffers generating neural representations. We use addition as g to match the alpha blending function, which clearly makes this function mathematically permutation invariant to inputs. This allows the neural network to blend all the transparency buffers into a single latent vector that is independent of the input order.

3.4 Rendering Model

We introduce our rendering model, *GlassNet* as shown in Fig. 2. *GlassNet* has four components: the scene encoder, \mathcal{F} ; the permutation invariant transparency buffer blending function, \mathcal{T} ; the final blending network, \mathcal{B} ; and the rendering network, \mathcal{R} . Both \mathcal{F} and \mathcal{B} use U-Net [32] as the network backend, whereas other networks use multi-layer CNN without down-sampling. All four components are trained jointly as one neural network. Full details of the model are shown in Fig. 3.

The scene encoding network, \mathcal{F} , accepts tensors containing the position, normal, texture, material parameters, ω_o , depth buffer, and direct lighting. It then generates a neural scene representation, σ . Similarly, each transparency buffer is processed by \mathcal{T} , generating the neural representation of transparent surfaces ρ . All the inputs use positional encoding

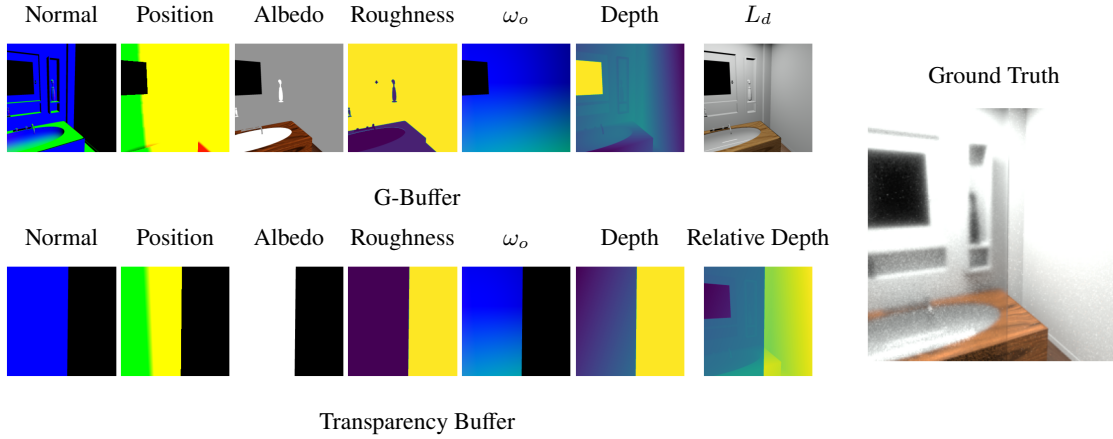


Fig. 5 G-buffer, transparency buffer, and ground truth. Our method can overcome the noise in the ground truth.

by Mildenhall *et al.* [33] to capture high-frequency information. Afterward, the blending network \mathcal{B} blends the scene representation σ and transparent object representations ρ and processes visibility. By doing so, \mathcal{B} generates the final neural buffer, ϕ .

Finally, the rendering network \mathcal{R} takes the neural buffer ϕ with additional direct lighting as input to render the final image. Although previous works, such as Neural Global Illumination, report using convolutional neural network structure on the rendering network can lead to distortion of shapes, we find that this problem can usually be solved by using kernels with the size of 1×1 at the end of the network, and adding larger kernel can help with global effects such as blurring. For balancing the render quality and performance, the *GlassNet* is trained with lower resolution (256×256) textures, G-buffers, and ground truths. A multi-layer CNN is used as a super-sampling network that super-samples the result at $\times 2$ the original resolution. Given that we can cheaply get high-resolution texture and G-buffer from the rasterizer, the super-sampling network can be simple to achieve real-time performance.

3.5 Training

We train the model in an end-to-end fashion to bake a specific scene by randomly sampling camera angles and positions within certain ranges of each scene following uniform distributions to render ground truth images. In particular, for each scene, we generate output image and G-buffer pairs for training. The ground truth is rendered by the modern path tracer, Mitsuba 3 [34]. For the G-buffer and direct lighting generation, we implemented our custom rasterization renderer using OpenGL in C++, with CUDA support to support neural network inference. As a loss function, we choose is the com-




Fig. 6 Randomly selected images from datasets. The transparent areas are highlighted by the blue masks in the second row.

bination of Structural Similarity Index Measure (SSIM) [35] and L_1 loss.

4 Experiments

4.1 Experimental Settings

The input and output resolution used for training are both 256×256 px. The super-sampling network uses 512×512 textures to super-sample the 256×256 *GlassNet* rendered results to 512×512 px, improving the visual quality. In addition, all inputs are positional encoded using the method proposed by Mildenhall *et al.*. All types of inputs and ground truth images are shown in Figure 5. We used a relatively low number of path tracing samples and did not pre-process the data with a denoiser.

We used Adam optimizer [36] to perform gradient descent with different hyperparameters on each scene. The dataset details and hyperparameters such as train/validation/test splits, learning rate and regularization for each scene can be found in Appendix A. The training time varies from 5 to 10 hours using 8 RTX 2080 Ti GPUs, depending on the scene.

As evaluation metrics, we use Mean Absolute Error (MAE), LPIPS [37], DSSIM [38], and Peak Signal-to-Noise Ratio




Fig. 7 Qualitative comparison. We compare with Active Exploration in several challenging scenes. Note that our method can preserve details of objects behind transparent surfaces with the synthesis of global illumination. All scenes are rendered on test sets with different camera views from the training set. The transparent areas are indicated by yellow dashed lines in the column of ground truths.

Scene	Method	Metrics					
		MAE ↓	LIPIPS ↓	DSSIM ↓	PSNR ↑	T.MAE ↓	T.PSNR ↑
BATHROOM	Active Exp.	0.01475	0.02995	0.04840	30.85680	0.06170	23.72751
	Neural G.I.	0.00934	0.00961	0.02280	35.03145	0.03888	28.31231
	Naive	0.02131	0.04409	0.05622	29.17893	0.08208	22.41524
	Ours	0.00808	0.00686	0.01887	36.81351	0.02738	32.46141
CORNELLBOX	Active Exp.	0.03237	0.18584	0.20583	25.21868	0.10491	19.78125
	Neural G.I.	0.04093	0.12171	0.19213	24.13943	0.13141	17.84513
	Naive	0.04988	0.20916	0.23296	22.90424	0.15471	17.80871
	Ours	0.02142	0.01595	0.05436	29.75453	0.06087	25.44725
DESK	Active Exp.	0.01739	0.05523	0.06117	27.87549	0.15792	16.31566
	Neural G.I.	0.01733	0.04327	0.04985	29.34936	0.12360	17.96721
	Naive	0.03274	0.07845	0.07987	26.16205	0.19205	15.75951
	Ours	0.01658	0.01150	0.01782	34.82217	0.05487	28.24681
CORNELLBOX8	Active Exp.	0.02911	0.17719	0.20583	26.16428	0.08924	21.14504
	Neural G.I.	0.04307	0.19579	0.27657	22.83434	0.13061	17.90722
	Naive	0.05093	0.28167	0.31939	21.84069	0.15147	17.09835
	Ours	0.02064	0.04152	0.07311	28.56370	0.05721	24.25974

Table 1 Quantitative comparison. We compare our approach with the naive baseline, Active Exploration and Neural Global Illumination on several scenes with diverse metrics. T.MAE and T.PSNR denote the MAE loss and PSNR on the areas of the image with transparent objects. Best result is denoted in **bold**.

(PSNR). For transparency quality evaluation, we use PSNR and MAE computed on pixels rendered with transparency (T.MAE and T.PSNR).

4.2 Comparison with the State-of-the-Art

We compared with Active Exploration and Neural Global Illumination, as they are the closest recent work. Comparisons were carried out in 4 challenging scenes: BATHROOM, CORNELLBOX, DESK, and CORNELLBOX8 as shown in Fig. 6. The BATHROOM was modified from the Bitterli dataset [39], while DESK and CORNELLBOX were created using existing models [40], to focus on transparency effects. In CORNELLBOX8, We added 4 more transparent objects to CORNELLBOX to explore the performance under the condition of high number of transparent objects. Qualitative results are shown in Fig. 1 and Fig. 7, while quantitative results are shown in Tab. 1. A comparison with real-time path tracing method is also demonstrated in Fig. 8

From the qualitative comparison in Fig. 7 we can see that although Active Exploration performed well in many areas, without transparency buffers, the neural network could not synthesize the area behind the semi-transparent surfaces without losing fundamental details. In the DESK and CORNELLBOX scenes, the text texture was almost impossible to render without provided texture information. Instead, our method could synthesize all the texture accurately. Furthermore, in the BATHROOM scene, all the detailed shapes behind the glass such as taps were missing in the Active Exploration, where our method preserved them well. The

Scene	Method	Metrics		
		DSSIM ↓	T.MAE ↓	T.PSNR ↑
CORNELLBOX	CNSR	0.1502	0.0973	20.4854
	Ours	0.0451	0.0457	26.2100
DESK	CNSR	0.0539	0.1115	18.5453
	Ours	0.0174	0.0533	28.6498

Table 2 Quantitative comparison with CNSR. Best result is highlighted in **bold**.

quantitative comparison in Tab. 1 corroborates the visual improvement of our method, which achieves better values on most metrics, with more significant improvements on parts of the image that need transparency. In addition, the rendering results and comparison video are available in Video S1 in the Electronic Supplementary Material.

Overall, our method preserves the details of texture and geometry of shape well. The global illumination effects such as mirror reflection are also well synthesized. Our approach shows strong performance on preserve all the high frequency texture, which is also shown on the wood texture of the BATHROOM scene. Such details are extremely hard to be synthesized behind nearly transparent surfaces without transparency buffers.

Our experiments show that the transparency buffer is essential for neural network to render objects behind transparent surfaces accurately. Although the global illumination effects are not the main focus of our method, we also demonstrate good quality on indirect diffuse lighting, mirror reflection, and soft shadows on par with existing approaches.

We also compared our method to CNSR to further verify the




Fig. 8 Comparison to real-time ray tracing. The real-time path tracing denoiser failed to reconstruct the high frequency texture and small objects behind the glass.

effectiveness of our approach to transparency. We generated the dataset in the format CNSR requires and used the identical dataset to train the CNSR model and our model. The results are shown in Fig. 9 and Tab. 2. The transparency quality rendered by our approach is significantly better. Note that this comparison experiment is only supplementary since the main focus of the two methods is different.

In addition, we compared our approach to the real-time path tracing method in Fig. 8. We used 32 samples per pixel to generate the path tracing image with a standard Optix denoiser [41]. The Optix denoiser failed to reconstruct the wood texture and the tape behind the glass, while our algorithm rendered such objects accurately. Due to the glass, the denoiser can no longer rely on albedo and normal AOV to gain extra information, causing the loss of details.

Size	Type	# of Trans. Obj.			
		1	2	4	8
256 × 256	Framerate (FPS)	100	77	63	45
	Frametime (ms)	10	13	16	22
512 × 512	Framerate (FPS)	50	37	32	23
	Frametime (ms)	20	27	31	43
1024 × 1024 [†] _{CARN}	Framerate (FPS)	24	21	20	18
	Frametime (ms)	41	47	50	55
1024 × 1024 [†] _{NSRR}	Framerate (FPS)	17	16	15	12
	Frametime (ms)	56	62	66	78

Table 3 Performance evaluation. We can interactively render all test cases. [†] 1024 × 1024 achieves by an external supersampling network, CARN or NSRR with temporal stability, as a post-processing procedure.

4.3 Ablation Study

Although various previous works such as Active Exploration and CNSR reported that the PixelGenerator has better up-sampling performance than CNN, we found that PixelGenerator is not able to effectively synthesize the transparent structure given the same amount of training time as CNN. We implemented the same rendering network structure with

	MAE↓	LIPIPS↓	DSSIM↓	PSNR↑
w/o P.E.	0.023711	0.026580	0.059718	28.937529
L1 Loss	0.023923	0.033820	0.072169	25.218679
Ours	0.021424	0.015953	0.054358	29.754528

Table 4 Impact of positional encoding and loss functions. Our experiments show the combination of positional encoding and L1SSIM loss function achieves the best performance of the network. Best result is highlighted in **bold**.

	MAE	MAE (inv. order)	Difference
BATHROOM	0.00805	0.00805	0.00000
DESK	0.01658	0.01658	0.00000
CORNELLBOX	0.02125	0.02125	0.00000

Table 5 MAE on different transparency buffer order. We invert the order of transparency buffer, and achieve identical rendering results.

PixelGenerator as the backbone network. Section 4.2 shows that PixelGenerator failed to reconstruct the glossy effects on transparent objects given 400 training iterations.

To overcome the bad up-sampling performance of CNN, an up-sampling network is solely trained to super-sample the rendered images at two times the resolution. Our experiment shows fine-tuning the model with a smaller high-resolution dataset cannot achieve good visual effects. Figure 13 shows that the image generated by our up-sampling network has better visual quality, especially on high-roughness transparent objects, than retraining the CNN.

We studied the influence of positional encoding on our method. Mildenhall *et al.* show that positional encoding can help neural networks learn high-frequency information; therefore, all the inputs to our neural network are positional encoded. We also demonstrate that using the combination of MAE loss and DSSIM loss can achieve the best quality, with quantitative results are shown in Tab. 4. Furthermore, we verified the effectiveness of our permutation-invariant algorithm. Table 5 shows that shuffled transparency buffer causes on effect on our model performance.

4.4 Performance Evaluation

4.4.1 Real-time Performance

We tested our renderer on one RTX 4090, and the performance is shown in Tab. 3. The computation time of our algorithm does not significantly grow w.r.t. the complexity of meshes because the transparent buffers only separately record each transparent object instead of each primitive. The CORNELLBOX and DESK (30,344 and 38,400 triangles of transparent objects) show that our method can efficiently render complex meshes.

Furthermore, the neural blending algorithm shown in Eq. (5) also grows linearly w.r.t. the number of transparent objects given a specific resolution because h is $n \mathcal{O}(1)$




Fig. 9 Rendering results of ours and Granskog *et al.* Our method preserves the transparency quality better. The transparent areas are indicated by yellow dashed lines.

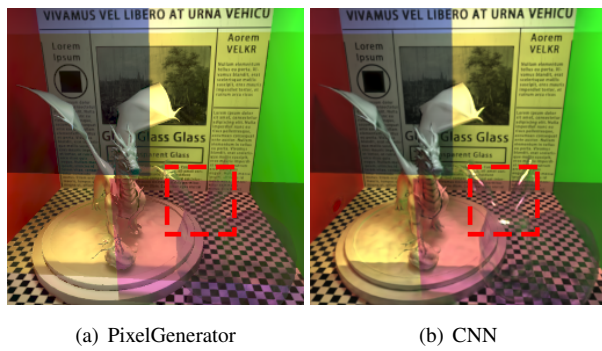


Fig. 10 PixelGenerator v.s. CNN. Under the same training iterations, PixelGenerator completely failed to capture the glossy effects on the glass bunny.

function under specific input size. With the resolution of 512×512 , one extra transparent object costs 3.3 ms more time on average. We believe this cost can be further optimized by pruning the model.

Our approach can be further supersampled by other existing fast supersampling networks. As examples, we demonstrate that by using CARN[42] and NSRR[43] as post-processing models, our method can reach 1024×1024 resolution with anti-aliasing effects. The results can be found in Tab. 3, Fig. 11, and Video S1 in the Electronic Supplementary Material.

4.4.2 Memory Efficiency

The permutation invariant blending algorithm also has a memory efficiency advantage over existing approaches. With traditional methods, the model has to receive all transparency buffers in order to make inference, so that the memory costs grow linearly w.r.t. the number of transparent objects. However, with our proposed algorithm, there is no additional framebuffer costs because the model can operate on a single buffer of each transparent objects, and then add it to the summed latent representation. The memory usage comparison




Fig. 11 Result in combining with other supersampling models. We demonstrate that using external supersampling models such as NSRR can further supersample our results to 1024×1024 .

	RAM (GB)	VRAM (GB)
Active Exploration	2.50	11
Ours	0.68	4

Table 6 Runtime memory usage comparison. We compare our method with Active Exploration on 512×512 with three transparent objects. Least memory usage is highlighted in **bold**.

is plotted in Fig. 12 approximated by framebuffer usage. We can see that our permutation invariant blending algorithm has constant space growth w.r.t. the number of transparent objects.

We also compare the actual runtime RAM and VRAM usage of our method with that of Active Exploration. The experiment uses the DESK scene with the resolution of 512×512 to test the memory performance. Table 6 shows our implementation significantly reduces the required runtime memory.

5 Limitations and Discussion

Although our method shows significant improvements in neural rendering of scenes with transparent objects with full OIT, our method relies on two fundamental assumptions about scenes: the presence of transparent surfaces with a low refraction index and the absence of participating media.




Fig. 12 Buffer memory usage comparison. Under 256×256 , one transparent object buffer causes around 4.5 megabytes of memory. Our blending algorithm is summation based and thus do not require space to store each transparency buffer.




Fig. 13 Fine-tuning v.s. Supersampling. Fine-tuning network on 512×512 data has bad quality of roughness effects. The window frame should be blurred out as in the ground truth.

However, when refraction occurs, the shading information provided by the G-buffer is no longer accurate, which becomes an even more significant issue on spherical surfaces. This limitation is illustrated in Fig. 14. Our rendering network currently struggles to produce correct results when strong refraction effects are present. Additionally, due to the use of G-buffers, our framework is unable to represent participating media, leading to the ineffective synthesis of volumetric effects.

As a problem of neural rendering in general, current neural rendering approaches perform relatively poorly on high-resolution images. Although our method uses CARN or NSRR as a post-processing supersampling procedure to boost the resolution further to 1024×1024 as described in Sec. 4.4.1, further attempts to raise resolution will significantly affect the performance, causing low framerates. Further research may focus on the performance of neural rendering in high-resolution settings.

Another potential direction for future research is incorporat-

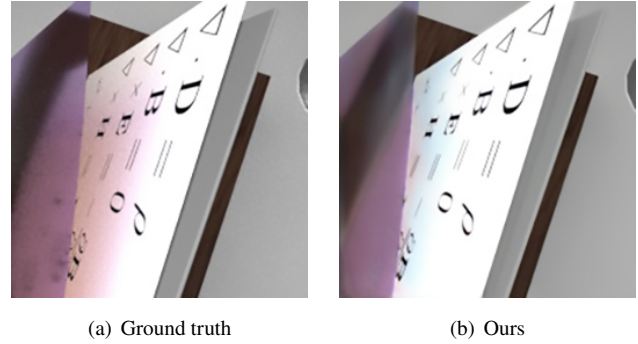


Fig. 14 Failed case. When the refraction effects dominate, our method can no longer accurately render surfaces under transparent objects because G-buffer cannot provide useful information.

ing path prediction techniques, such as the method proposed by Li *et al.* [44], into scene representation. Furthermore, exploring efficient approaches for representing participating media on both the input and neural network sides presents an intriguing research topic.

Several orthogonal methods can be combined with our work. For example, the Monte Carlo sampling technique proposed in Active Exploration can be utilized to improve the effects such as perfect mirror reflection and caustic effects. Neural Global Illumination uses radiance cues as inputs to the model, which improves the scene representation learning.

Appendix

Table 7 records the key scene settings and hyperparameters during the dataset generation, training, and evaluation.

Declaration of competing interest

The authors have no competing interests to declare that are relevant to the content of this article. No funding was received to assist with the preparation of this manuscript.

Electronic Supplementary Material

The rendering video recorded in real-time is available in the electronic supplementary material of this article. The code is publicly available on <https://github.com/ziyangz5/NeuralBakingTransparency/>.

References

- [1] Kajiya JT. The rendering equation. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, 1986, 143–150. 1, 4
- [2] Lafortune EP, Willems YD. Rendering Participating Media with Bidirectional Path Tracing. In X Pueyo, P Schröder, editors, *Rendering Techniques '96*, 1996, 91–100. 1

Scene	Train	Validation	Test	S.P.P.	L.R.	W.D.	# of Trans. Obj.	Variable
BATHROOM	1800	200	128	8192	0.0001	0.00075	1	1 (Roughness)
CORNELLBOX	1800	200	128	4096	0.0001	0.0001	4	4 (Position and Lighting)
DESK	1800	200	128	4096	0.0001	0.0001	3	3 (Color)
CORNELLBOX8	900	100	128	4096	0.0001	0.0001	8	7 (Position and Lighting)

Table 7 Scene settings and hyperparameter choices. S.P.P. indicates the number of samples per pixel using during the ground truth generation. L.R. refers to learning rate and W.D. refers to weight decay.

- [3] Greger G, Shirley P, Hubbard PM, Greenberg DP. The irradiance volume. *IEEE Computer Graphics and Applications*, 1998, 18(2): 32–43. 1, 3
- [4] Ren P, Wang J, Gong M, Lin S, Tong X, Guo B. Global illumination with radiance regression functions. *ACM Trans. Graph.*, 2013, 32(4): 130–1. 2, 4
- [5] Eslami SA, Jimenez Rezende D, Besse F, Viola F, Morcos AS, Garnelo M, Ruderman A, Rusu AA, Danihelka I, Gregor K, et al.. Neural scene representation and rendering. *Science*, 2018, 360(6394): 1204–1210. 2, 4
- [6] Granskog J, Rousselle F, Papas M, Novák J. Compositional neural scene representations for shading inference. *ACM Transactions on Graphics (TOG)*, 2020, 39(4): 135–1. 3, 4
- [7] Rainer G, Bousseau A, Ritschel T, Drettakis G. Neural Pre-computed Radiance Transfer. *Computer Graphics Forum (Proceedings of the Eurographics conference)*, 2022, 41(2).
- [8] Diolatzis S, Philip J, Drettakis G. Active exploration for neural global illumination of variable scenes. *ACM Transactions on Graphics (TOG)*, 2022, 41(5): 1–18. 2, 3
- [9] Gao D, Mu H, Xu K. Neural Global Illumination: Interactive Indirect Illumination Prediction under Dynamic Area Lights. *IEEE Transactions on Visualization and Computer Graphics*, 2022. 2, 3, 4
- [10] Porter T, Duff T. Compositing digital images. In *Proceedings of the 11th annual conference on Computer graphics and interactive techniques*, 1984, 253–259. 2, 5
- [11] Carpenter L. The A-buffer, an antialiased hidden surface method. In *Proceedings of the 11th annual conference on Computer graphics and interactive techniques*, 1984, 103–108. 2
- [12] Jouppi NP, Chang CF. Z 3: an economical hardware technique for high-quality antialiasing and transparency. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS workshop on Graphics hardware*, 1999, 85–93. 2
- [13] Everitt C. Interactive order-independent transparency. *White paper, nVIDIA*, 2001, 2(6): 7. 3
- [14] Saito T, Takahashi T. Comprehensible rendering of 3-D shapes. In *Proceedings of the 17th annual conference on Computer graphics and interactive techniques*, 1990, 197–206. 3
- [15] Pangerl D. Deferred rendering transparency. *ShaderX7: Advanced Rendering Techniques, ShaderX series*, 2009: 217–224. 3
- [16] Mara M, McGuire M, Luebke D. Lighting deep g-buffers: Single-pass, layered depth images with minimum separation applied to indirect illumination. *NVIDIA Corporation*, 2013. 3
- [17] Xin H, Zheng S, Xu K, Yan LQ. Lightweight bilateral convolutional neural networks for interactive single-bounce diffuse indirect illumination. *IEEE Transactions on Visualization and Computer Graphics*, 2020, 28(4): 1824–1834. 3, 4
- [18] Keller A, Fascione L, Fajardo M, Georgiev I, Christensen P, Hanika J, Eisenacher C, Nichols G. The Path Tracing Revolution in the Movie Industry. In *ACM SIGGRAPH 2015 Courses*, 2015, 1–7, doi:10.1145/2776880.2792699. 3
- [19] Ritschel T, Dachsbacher C, Grosch T, Kautz J. The state of the art in interactive global illumination. In *Computer graphics forum*, volume 31, 2012, 160–188. 3
- [20] O'Donnell Y. Precomputed Global Illumination in Frostbite, 2018. 3
- [21] Cohen MF, Wallace JR, Hanrahan P. Radiosity and realistic image synthesis, 1993. 3
- [22] McGuire M, Mara M, Nowrouzezahrai 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, 1–11. 3
- [23] Tewari A, Fried O, Thies J, Sitzmann V, Lombardi S, Sunkavalli K, Martin-Brualla R, Simon T, Saragih J, Nießner M, et al.. State of the art on neural rendering. In *Computer Graphics Forum*, volume 39, 2020, 701–727. 4
- [24] Raghavan N, Xiao Y, Lin KE, Sun T, Bi S, Xu Z, Li TM, Ramamoorthi R. Neural Free-Viewpoint Relighting for Glossy Indirect Illumination. In *Computer Graphics Forum*, volume 42, 2023, e14885. 4
- [25] Qi CR, Su H, Mo K, Guibas LJ. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, 652–660. 4, 5
- [26] Bloem-Reddy B, Teh YW. Probabilistic symmetries and invariant neural networks. *The Journal of Machine Learning Research*, 2020, 21(1): 3535–3595. 4
- [27] Heitz E, Dupuy J, Hill S, Neubelt D. Real-time polygonal-light shading with linearly transformed cosines. *ACM Transactions on Graphics (TOG)*, 2016, 35(4): 1–8. 4
- [28] Wloka M. Batch, batch, batch: What does it really mean, 2003. 5
- [29] Pantazopoulos I, Tzafestas S. Occlusion culling algorithms: A comprehensive survey. *Journal of Intelligent and Robotic Systems*, 2002, 35: 123–156. 5
- [30] Meshkin H. Sort-independent alpha blending. *GDC Talk*, 2007, 2(4). 5

- [31] Salvi M, Vaidyanathan K. Multi-layer alpha blending. In *Proceedings of the 18th meeting of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, 2014, 151–158. 5
- [32] Isola P, Zhu JY, Zhou T, Efros AA. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, 1125–1134. 5
- [33] Mildenhall B, Srinivasan PP, Tancik M, Barron JT, Ramamoorthi R, Ng R. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 2021, 65(1): 99–106. 6
- [34] Jakob W, Speierer S, Roussel N, Vicini D. Dr.Jit: A Just-In-Time Compiler for Differentiable Rendering. *Transactions on Graphics (Proceedings of SIGGRAPH)*, 2022, 41(4), doi: 10.1145/3528223.3530099. 6
- [35] Wang Z, Bovik AC, Sheikh HR, Simoncelli EP. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 2004, 13(4): 600–612. 6
- [36] Kingma DP, Ba J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [37] Zhang R, Isola P, Efros AA, Shechtman E, Wang O. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, 586–595. 6
- [38] Loza A, Mihaylova L, Canagarajah N, Bull D. Structural similarity-based object tracking in video sequences. In *2006 9th International Conference on Information Fusion*, 2006, 1–6. 6
- [39] Bitterli B. Rendering resources, 2016, <https://benedikt-bitterli.me/resources/>. 8
- [40] Zhou K, Wang X, Tong Y, Desbrun M, Guo B, Shum HY. Texturemontage: Seamless Texturing of Arbitrary Surfaces From Multiple Images. *ACM Transactions on Graphics*, 2005, 24(3): 1148–1155. 8
- [41] Chaitanya CRA, Kaplanyan AS, Schied C, Salvi M, Lefohn A, Nowrouzezahrai D, Aila T. Interactive reconstruction of Monte Carlo image sequences using a recurrent denoising autoencoder. *ACM Transactions on Graphics (TOG)*, 2017, 36(4): 1–12. 9
- [42] Ahn N, Kang B, Sohn KA. Fast, accurate, and lightweight super-resolution with cascading residual network. In *Proceedings of the European conference on computer vision (ECCV)*, 2018, 252–268. 10
- [43] Xiao L, Nouri S, Chapman M, Fix A, Lanman D, Kaplanyan A. Neural supersampling for real-time rendering. *ACM Transactions on Graphics (TOG)*, 2020, 39(4): 142–1. 10
- [44] Li Z, Yeh YY, Chandraker M. Through the looking glass: Neural 3d reconstruction of transparent shapes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, 1262–1271. 11

Author biography



Ziyang Zhang is currently a Ph.D. candidate in Faculty of Science and Engineering at Waseda University. He received his bachelor of science degree in computer science from University of California, Irvine in 2020, and his master of science degree in computer science from University of California, San Diego in 2022. His research interests include photorealistic rendering, real-time rendering and machine learning with computer graphics.



Edgar Simo-Serra is currently an associate professor at Waseda University. He obtained his Industrial Engineering degree from BarcelonaTech in 2011 and his Ph.D. in 2015 from the same university. From 2015 to 2018 he was at Waseda University as a junior researcher (assistant professor), and during 2018 he was a JST Presto Researcher, before re-joining Waseda. His general research interests are in the intersection of computer vision, computer graphics, and machine learning with applications to large-scale real world problems.