

RenderFlow: Single-Step Neural Rendering via Flow Matching

Shenghao Zhang^{1,2} Runtao Liu¹ Christopher Schroers¹ Yang Zhang¹

¹Disney Research|Studios ²ETH Zürich

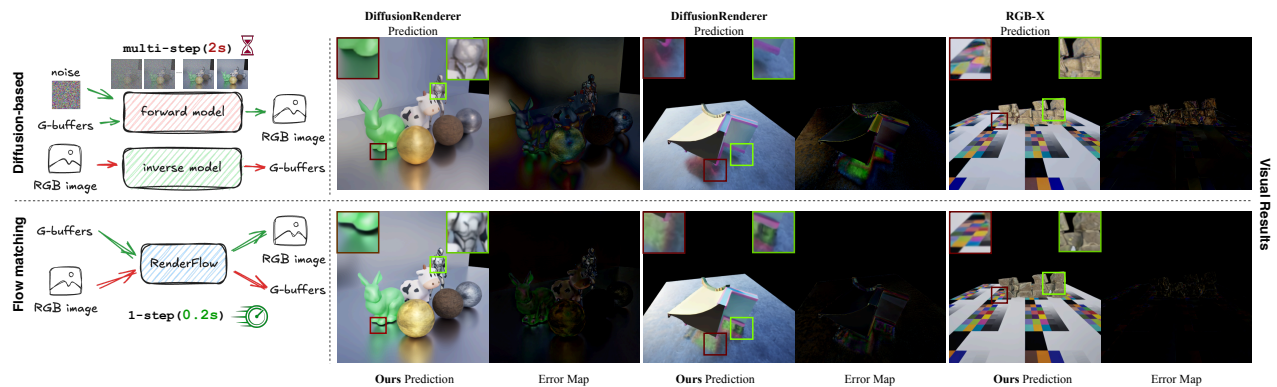


Figure 1. **Overview of the proposed *RenderFlow* framework.** **Left:** *RenderFlow* learns a single-step conditional flow from *albedo* (rather than noise) to shaded images, yielding $\sim 10\times$ faster rendering while better physical light transport; both forward and inverse rendering are unified within a *single* model. **Right:** Example visualizations and corresponding error maps illustrating fidelity.

Abstract

Conventional physically based rendering (PBR) pipelines generate photorealistic images through computationally intensive light transport simulations. Although recent deep learning approaches leverage diffusion model priors with geometry buffers (G-buffers) to produce visually compelling results without explicit scene geometry or light simulation, they remain constrained by two major limitations. First, the iterative nature of the diffusion process introduces substantial latency. Second, the inherent stochasticity of these generative models compromises physical accuracy and temporal consistency. In response to these challenges, we propose a novel, end-to-end, deterministic, single-step neural rendering framework, *RenderFlow*, built upon a flow matching paradigm. To further strengthen both rendering quality and generalization, we propose an efficient and effective module for sparse keyframe guidance. Our method significantly accelerates the rendering process and, by optionally incorporating sparsely rendered keyframes as guidance, enhances both the physical plausibility and overall visual quality of the output. The resulting pipeline achieves near real-time performance with photorealistic rendering quality, effectively bridging the gap between the efficiency of modern

generative models and the precision of traditional physically based rendering. Furthermore, we demonstrate the versatility of our framework by introducing a lightweight, adapter-based module that efficiently repurposes the pre-trained forward model for the inverse rendering task of intrinsic decomposition.

1. Introduction

The pursuit of photorealism in computer graphics has long been driven by physically based rendering (PBR) techniques, which simulate the complex interaction of light, materials, and geometry to achieve high visual fidelity [10, 21]. Methods such as path tracing are the gold standard for off-line applications, such as visual effects in film productions, but their high computational cost, requiring extensive sampling to obtain “noise-free” images, makes them impractical for real-time settings like video games, interactive virtual production, and pre-visualization, where high frame rates and rapid creative iteration are essential.

Recent progress in large-scale generative models, especially latent diffusion models [25], have enabled neural rendering methods that synthesize photorealistic images from

G-buffer inputs [16, 37]. Despite their strong visual fidelity, these approaches face key limitations: the iterative denoising process typically requires 20–50 network evaluations, creating latency unsuitable for interactive use; and the stochastic sampling inherent to diffusion models reduces physical accuracy and temporal stability. As a result, flickering and other inconsistencies often emerge, preventing these methods from meeting studio-quality standards and limiting their practical deployment.

To bridge the gap between the generative power of diffusion models and the demands of real-time, physically-consistent rendering, we propose *RenderFlow*. Our method reformulates neural rendering as a single-step conditional flow generation process. The core idea is to leverage a flow matching paradigm [17] to learn a velocity flow field from the diffuse albedo, rather than Gaussian noise, to the final fully-shaded image, conditioned on the G-buffer attributes (e.g., normals, depth, materials). By replacing noise with albedo as shown in Figure 1, we preserve low-level signals while the network synthesizes high-frequency details, maintaining the geometric integrity of the scene. Learning this direct residual flow enables *RenderFlow* to leverage pre-trained generative priors for synthesizing complex shading and lighting effects in a single forward pass.

However, achieving physically accurate results remains challenging for neural rendering methods, which lack explicit representations of scene geometry and light transport. To address this and improve both physical accuracy and temporal stability, we introduce a novel guidance mechanism that leverages sparse, high-quality reference frames rendered offline via path tracing. These frames provide strong conditioning that anchors the generative process, ensuring the output remains faithful to ground-truth light transport. This optional module effectively reduces visual artifacts and improves consistency, making the method more robust for practical applications.

Furthermore, to demonstrate the versatility of our framework, we introduce a parameter-efficient adaptation for the inverse rendering task. By augmenting the frozen *RenderFlow* backbone with adapter modules, our model can be efficiently repurposed to decompose a rendered image back into its intrinsic G-buffer components. This addition establishes our approach as a unified paradigm, capable of handling both forward (synthesis) and inverse (decomposition) rendering within a single, shared architecture, which can be potentially applied in further video editing applications.

Our main contributions are:

- A single-step, flow-based rendering framework, *RenderFlow*, that achieves a faster inference speed by reformulating neural rendering as a direct conditional flow generation task, significantly outperforming iterative diffusion-based methods in quality and efficiency.
- A novel reference-frame guidance module that improves

the physical accuracy of the rendered output, achieving superior visual quality and mitigating artifacts common in prior generative approaches.

- A unified, efficient adapter-based inverse renderer built on a frozen forward renderer, demonstrating the versatility and extensibility of *RenderFlow*.

2. Related Work

2.1. Neural Rendering

Early works employed Convolutional Neural Networks (CNNs) [19] and Generative Adversarial Networks (GANs) [30] to approximate screen-space shading from G-buffers. More recently, Transformer-based architectures like Lightformer [24] and Renderformer [36] have been explored to embed lighting and mesh properties explicitly. The latest generation of methods leverages large-scale diffusion models for both inverse rendering and forward rendering conditioned on G-buffer attributes [16, 37], with some works exploring unified models for joint albedo estimation and relighting [7]. While these approaches produce visually compelling results, they are constrained by two critical limitations: the iterative, multi-step sampling process introduces prohibitive latency for real-time applications, and the inherent stochasticity of the generative process compromises the physical fidelity required for production-quality rendering.

2.2. Efficient Generative Models

To overcome the latency of iterative sampling, significant research has focused on accelerating generative models. One major line of work aims to distill pre-trained, multi-step diffusion models into a single inference step. Foundational techniques such as adversarial diffusion distillation [27] and distribution matching distillation [35] have proven effective at training fast, one-step generators that maintain high fidelity. These distilled models can be efficiently fine-tuned for various downstream image-to-image translation tasks such as face restoration, superresolution, and 3D reconstruction enhancement [20, 32, 33, 38].

An alternative and complementary approach is offered by **Flow Matching** models, such as Rectified Flow [17]. Flow matching learns a velocity vector field that transports samples from a simple prior distribution such as a Gaussian distribution to the target data distribution by solving an ordinary differential equation (ODE). This formulation is highly efficient, as the ODE can be solved with very few function evaluations. Large-scale flow matching models have demonstrated state-of-the-art performance in image and video synthesis [6, 15, 31]. Recent works have demonstrated the versatility of flow matching in various downstream tasks including image relighting, shadow generation [5]. Our work builds upon these advancements, proposing a one-step, flow-based framework specifically tailored for

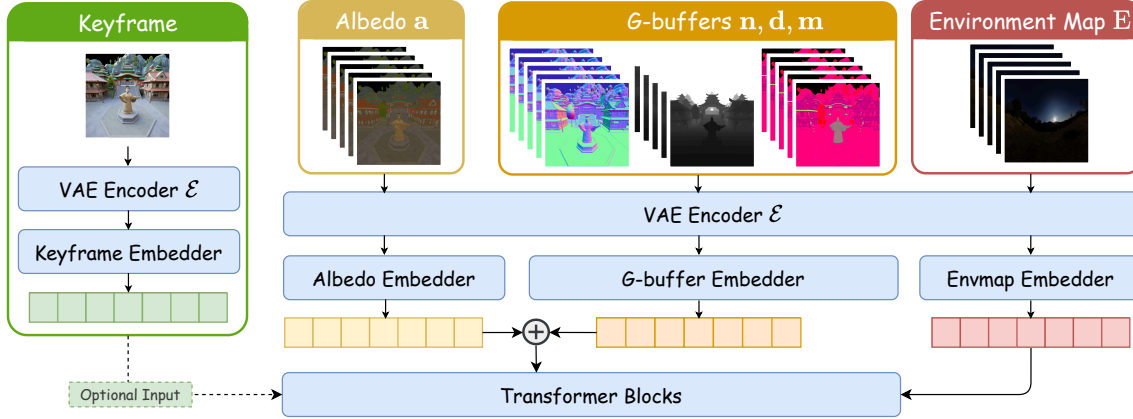


Figure 2. **An overview of our proposed RenderFlow architecture.** Unlike diffusion methods that start from noise, RenderFlow takes albedo as input and directly predicts fully-shaded outputs. Built on a pre-trained video DiT, it injects aligned G-buffer tokens into transformer blocks. Environment maps and optional keyframes are integrated via lightweight adapters (see Fig. 3 for details). The transformer aggregates all inputs to produce temporally coherent, physically accurate rendering outputs.

efficient and high-fidelity neural rendering.

3. Preliminaries

3.1. Classical Rendering Frameworks

Physically-based rendering (PBR) pipelines approximate this integral with Monte Carlo path tracing. Although path tracing delivers photorealistic images by simulating light transport, its heavy computational rules it out for interactive contexts. Deferred rendering attains real-time performance by decoupling geometry from shading, where surface attributes (e.g., normals, albedo) are stored in screen-space G-buffers. A subsequent lighting pass then uses these buffers for shading, though often at the cost of the global illumination accuracy provided by path tracing.

3.2. Diffusion and Bridge Matching Models

Denosing Diffusion Probabilistic Models (DDPM) [8] define a forward Markov process that gradually adds Gaussian noise to data \mathbf{x}_0 , transforming it into a prior distribution. The reverse process learns to iteratively denoise a sample from the prior back into the data distribution by training a network ϵ_θ to predict the added noise at each timestep t .

A limitation of standard diffusion models is their dependence on a fixed Gaussian prior. The **bridge matching** framework [2, 28] generalizes this by learning a stochastic trajectory between two arbitrary distributions π_0 and π_1 . Given samples $\mathbf{x}_0 \sim \pi_0$ and $\mathbf{x}_1 \sim \pi_1$, the trajectory is defined by a stochastic differential equation (SDE):

$$d\mathbf{x}_t = v_\theta(\mathbf{x}_t, t)dt + \sigma(t)d\mathbf{W}_t, \quad (1)$$

where v_θ is the learnable velocity (drift) field and $\sigma(t)d\mathbf{W}_t$ adds stochastic perturbation via a time-dependent noise schedule, $\sigma(t)$. A common choice, the *Brownian bridge*,

uses $\sigma(t) = \sigma\sqrt{t(1-t)}$ for smooth interpolation. The velocity field is trained with:

$$\mathcal{L}_{\text{Bridge}} = \mathbb{E}_{\mathbf{x}_0, \mathbf{x}_1, t} \left[\left\| v_\theta(\mathbf{x}_t, t) - \frac{\mathbf{x}_1 - \mathbf{x}_t}{1-t} \right\|^2 \right], \quad (2)$$

In the deterministic case where $\sigma(t) = 0$, this reduces to **flow matching**, which learns a straight-line ODE path:

$$\mathcal{L}_{\text{Flow}} = \mathbb{E}_{\mathbf{x}_0, \mathbf{x}_1, t} [\|v_\theta(\mathbf{x}_t, t) - (\mathbf{x}_1 - \mathbf{x}_0)\|^2], \quad (3)$$

where $\mathbf{x}_t = (1-t)\mathbf{x}_0 + t\mathbf{x}_1$.

4. Method

We formulate neural rendering as a conditional flow-based generative modeling problem, learning a transformation from structured G-buffer attributes to the distribution of photorealistic rendered sequences. Leveraging the powerful prior of a large-scale pretrained video diffusion model, we reinterpret its dynamics as a conditional latent flow. We introduce a novel architecture and a training scheme based on the principles of bridge matching [17], enabling high-fidelity rendering in a single inference step.

4.1. Input Representation

Following the deferred rendering paradigm, our model processes a set of G-buffers. The primary input is the **Albedo** buffer $\mathbf{a} \in \mathbb{R}^{H \times W \times 3}$, which provides the base color. Geometric information is supplied by the **Normal** buffer $\mathbf{n} \in \mathbb{R}^{H \times W \times 3}$ (in camera space) and a single-channel **Depth** buffer $\mathbf{d} \in \mathbb{R}^{H \times W}$. Material properties, adhering to the widely adopted Disney BRDF principle [3, 14], are described by three single-channel buffers: roughness, metallic, and specular. These are concatenated along the channel axis to form a single three-channel **Material** buffer \mathbf{m} .

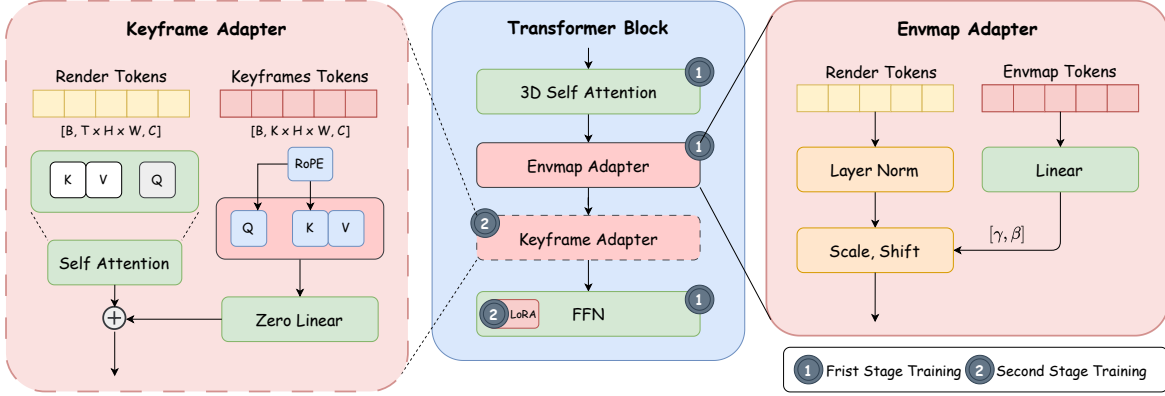


Figure 3. **Design of the customized transformer block in RenderFlow.** To incorporate global lighting context, the environment map is injected via an adaptive normalization layer before the transformer layers. To compensate for the limited information in per-frame G-buffers, the keyframe adapter introduces cross-attention to extract complementary temporal features from an *optional* keyframe. The adapter includes a LoRA module for parameter-efficient adaptation and employs Rotary Position Embeddings (RoPE) to encode temporal distances between keyframes and current frames.

For lighting conditions, we follow prior work [13, 16] and represent illumination with a high-dynamic-range (HDR) environment map $\mathbf{E} \in \mathbb{R}^{H' \times W' \times 3}$. By contrast, we simplify the input representation by rotating the environment map into the camera’s view space and applying Reinhard tone mapping operator [23] to produce a single LDR image, \mathbf{E}_{ldr} . We hypothesize that, at high resolution, a tone mapped environment map is sufficient for capturing the lighting information, making extra intensity channels in other dynamic ranges unnecessary. Rotating the map to the camera view during training allows the network to model the directional lighting implicitly, without explicit directional encodings.

4.2. Model Architecture

Backbone. Our architecture is built upon a pre-trained video diffusion model, Wan2.1 [31], which features a 3D causal VAE and a Diffusion Transformer (DiT) backbone. As illustrated in Fig. 2, all G-buffer attributes and the tone mapped environment map are first processed by the VAE encoder to obtain latent representations: $\mathcal{E}(\mathbf{a}), \mathcal{E}(\mathbf{n}), \mathcal{E}(\mathbf{d}), \mathcal{E}(\mathbf{m}), \mathcal{E}(\mathbf{E}_{\text{ldr}})$. To repurpose the model for our task, we replace the standard noisy latent input with a clean albedo latent $\mathcal{E}(\mathbf{a})$. An input embedder then patches the albedo latent into render tokens. Following VACE [12], the remaining spatially aligned G-buffer latents are processed by a decoupled attribute embedder. Given their spatial alignment, the resulting attribute tokens are added element-wise to the render tokens before being fed into the transformer blocks.

Envmap Adapter. Unlike the G-buffers, the environment map \mathbf{E}_{ldr} lacks spatial alignment with the render tokens. Instead, it functions as a global condition, governing the overall color palette and lighting characteristics of the

scene. To inject this global information, we modulate the render tokens within each transformer block with an adaptive normalization layer [34]. Specifically, the latent representation of the environment map is processed into envmap tokens. These tokens are then passed through a linear projection to predict the scale factor (γ) and shift (β) parameter. The projected render features are defined as the linear projection $f_{i+1} = (\gamma + 1) * f_i + \beta$, where f_i represents the render features at layer i .

Keyframe Guidance. To enable optional sparse-keyframe guidance for improved physical accuracy and visual quality, we introduce a Keyframe Adapter module (Fig. 3). This module adds a cross-attention branch in parallel to the existing self-attention layer, whose output is added to the original feature as a residual correction term. To encode the temporal distance between the current frame and each keyframe, we apply Rotary Position Embeddings (RoPE) [29] to the key and query vectors. Since this adapter is an optional component, we adopt a two-stage training strategy: Stage 1 trains the base model alone to learn the core rendering task, and Stage 2 freezes it while training only the Keyframe Adapter. This ensures that baseline performance remains unchanged when no keyframes are provided.

4.3. Latent-Space Bridge-Matching Training

We train our model under the stochastic bridge matching framework [2], aiming to learn a conditional trajectory that transforms the input albedo distribution to the distribution of fully rendered image. For computational efficiency, we formulate this process in the latent space. Let $\mathbf{z}_0 = \mathcal{E}(\mathbf{a})$ be the latent of the input albedo and $\mathbf{z}_1 = \mathcal{E}(\mathbf{I}_{\text{gt}})$ be the latent of the ground truth image. During training, we sample a timestep $t \sim \pi(t)$ where $t \in [0, 1]$ and interpolate be-

tween the start and end points of the path using the Brownian bridge formulation:

$$\mathbf{z}_t = (1-t)\mathbf{z}_0 + t\mathbf{z}_1 + \sigma\sqrt{t(1-t)}\epsilon \quad (4)$$

where $\epsilon \sim \mathcal{N}(0, \mathbf{I})$. The network takes the interpolated latent \mathbf{z}_t , the timestep t , and conditioning information as input to predict the conditional v_θ . The model is trained to minimize the bridge matching loss in the latent space:

$$\mathcal{L}_{\text{latent}} = \mathbb{E}_{\mathbf{z}_0, \mathbf{z}_1, t} \left[\left\| v_\theta(\mathbf{z}_t, t) - \frac{\mathbf{z}_1 - \mathbf{z}_t}{1-t} \right\|^2 \right] \quad (5)$$

Crucially, with a trained velocity field, v_θ , the final latent $\hat{\mathbf{z}}_1$ can be derived directly from the interpolated latent \mathbf{z}_t :

$$\hat{\mathbf{z}}_1 = \mathbf{z}_t + v_\theta(\mathbf{z}_t, t)(1-t). \quad (6)$$

Inspired by image translation literature [5, 20], we decode the predicted latent $\hat{\mathbf{z}}_1$ to the final rendered image $\mathbf{I}_{\text{pred}} = \mathcal{D}(\hat{\mathbf{z}}_1)$ and apply additional pixel-space losses to accelerate convergence and enhance reconstruction fidelity. Specifically, we optimize a composite pixel-wise loss:

$$\mathcal{L}_{\text{pixel}} = \mathcal{L}_{\text{LPIPS}} + \mathcal{L}_{\text{grad}}, \quad (7)$$

which combines the LPIPS loss [39] for perceptual similarity and a gradient loss [18] to reconstruct high-frequency details like contact shadows. The final training objective is a weighted sum of the latent and pixel losses:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{latent}} + \lambda\mathcal{L}_{\text{pixel}}. \quad (8)$$

4.4. Long video inference

Using pixel-space losses requires decoding latent representations at every training step, which is computationally expensive for long video sequences. We therefore train the model on short, fixed-length clips (typically 5 frames) and employ a progressive inference strategy for long videos. Specifically, during training, we randomly condition the model on the first frame of the input sequence and predict subsequent frames by feeding a masked reference clip, following VACE [12]. At inference time, we generate videos in overlapping chunks, where the last frame of the previous chunk serves as the conditioning frame for the next, promoting smooth transitions and temporal coherence over the full sequence.

4.5. Inverse Rendering

Our framework can be effectively repurposed for inverse rendering tasks without training a separate model from scratch. Concretely, we freeze the pretrained forward backbone and introduce a lightweight inverse adapter. An input RGB image is first encoded by the frozen VAE encoder \mathcal{E}

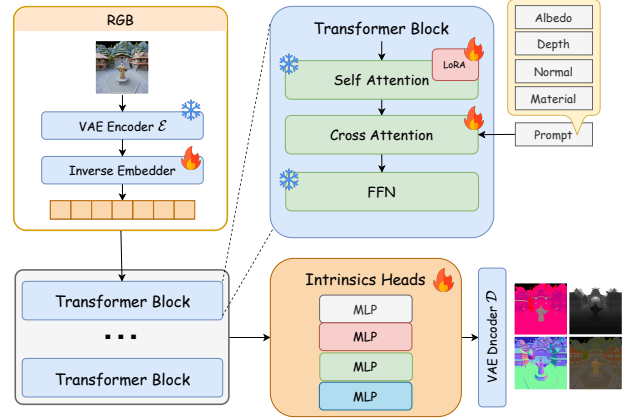


Figure 4. **Inverse adapter architecture.** The inverse adapter uses the frozen forward rendering backbone and introduces specialized modules to adapt it for G-buffer decomposition.

into a latent z_{rgb} , which is then mapped to tokens by a trainable inverse embedder. To repurpose the frozen transformer for intrinsic decomposition, we augment its self-attention projections with low-rank adaptation (LoRA) modules and insert a prompt-conditioned cross-attention branch where a text prompt selects the target intrinsic to be predicted, following the control mechanism popularized in RGB-X [37]. After the final transformer block, a set of lightweight, per-intrinsic heads (MLPs) project the tokens into the corresponding latent space, which is then decoded by the frozen VAE decoder \mathcal{D} to produce the requested G-buffer. This design preserves the representational rendering features of the forward model, enables parameter-efficient adaptation, and unifies multiple inverse tasks within a single model switchable by prompt, as illustrated in Fig. 4.

Training optimizes only the adapter parameters (inverse embedder, LoRA, cross-attention, and per-intrinsic heads) using modality specific reconstruction objectives: L1+LPIPS for albedo, cosine similarity for normals, scale- and shift-invariant loss for depth [22], and L1 for material components. In practice, this yields accurate and stable decomposition while retaining the forward renderer unchanged and avoiding full model retraining.

4.6. Dataset Synthesis

Due to the lack of a large-scale rendering dataset that contains the buffers and environment maps required in our task, we constructed a large-scale dataset using Unreal Engine 5 to train our model. Leveraging the *Movie Render Queue*, we rendered high-quality, path-traced reference images alongside their corresponding G-buffers. The dataset is composed of two distinct categories: large-scale, artist-crafted environments, and procedurally generated synthetic scenes. To overcome the computational expense and limited availability of complex, pre-existing scenes, we developed a pro-

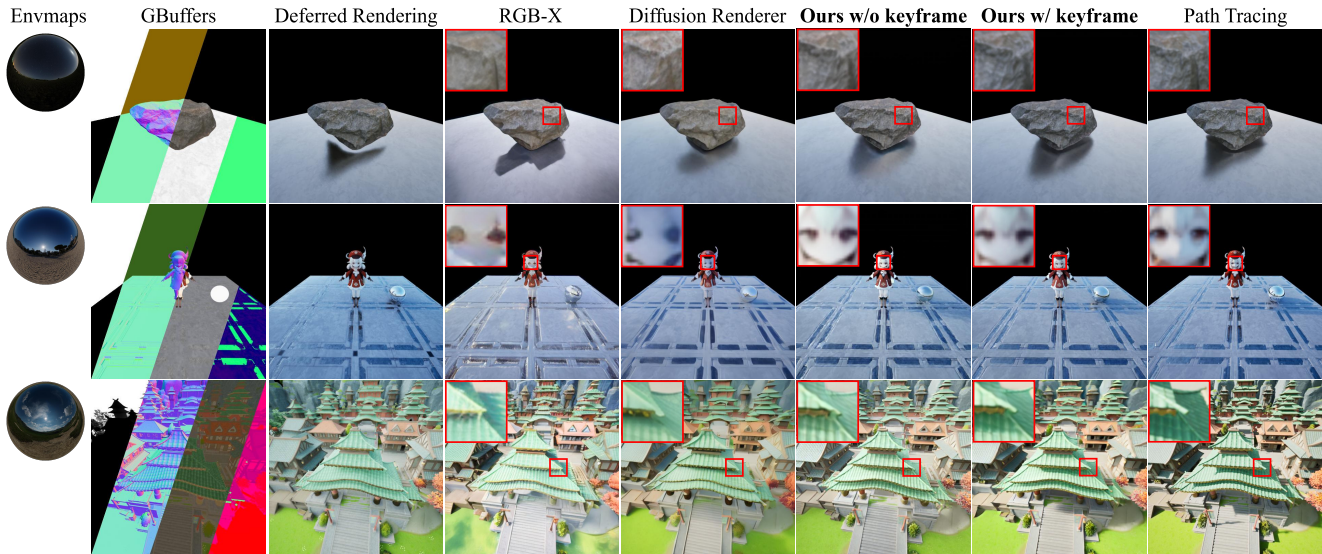


Figure 5. **Qualitative comparison of rendered images with traditional methods including deferred rendering, and diffusion-based approaches (RGB-X and DiffusionRenderer).** Our method better preserves texture details and produces high-quality shadows and lighting effects.

cedural pipeline to compose novel scenes from a curated collection of assets, significantly increasing data diversity. This collection includes approximately 4,000 unique meshes and 30 high-dynamic-range imaging (HDRI) environment maps. To further enrich the dataset, we randomly configured the material attributes of the meshes. In total, our dataset contains over 30,000 frames from artist-crafted environments and 100,000 frames from synthetic scenes. All images were rendered at a 512x512 resolution with 256 samples per pixel (SPP) and subsequently denoised using Intel’s Open Image Denoise [1].

5. Experiments

We refer to implementation details in the Supplementary.

5.1. Evaluation

Baselines. We compare against two state-of-the-art diffusion-based neural rendering frameworks: RGB-X [37] and DiffusionRenderer [16]. To ensure a fair comparison and mitigate potential domain gaps, we fine-tuned both baselines on our own dataset, following the training setups in their original papers. RGB-X was fine-tuned for 100,000 steps, while DiffusionRenderer, which operates on video frames, was fine-tuned on 4-frame sequences for 25,000 steps. At inference, we use 50 sampling steps for RGB-X and 30 steps for DiffusionRenderer.

Test Set and Metrics. We curated a challenging test set comprising both synthetic and artist-crafted scenes featuring complex lighting and shadows. The set includes 1,000 paired frames for synthetic scenes and 1,000 frames for

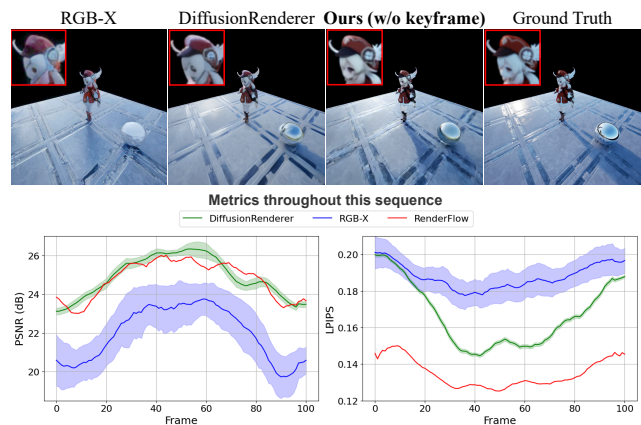


Figure 6. **Metrics are evaluated over a 100-frame sequence and averaged across 10 runs.** Shaded regions indicate variance across runs. Our method not only achieves the best overall performance but also zero variance, reflecting its deterministic, consistently high-quality behavior in contrast to the stochastic baselines.

artist-crafted scenes. We evaluate all methods using standard image reconstruction metrics: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), and the Learned Perceptual Image Patch Similarity (LPIPS) [39].

Quantitative Results. As shown in Tab. 1, our single-step method achieves superior performance to existing diffusion-based methods. Without keyframe guidance, it outperforms both RGB-X and DiffusionRenderer across all metrics. When keyframe guidance is enabled, the full model attains the best overall performance, with substantial improvements on every metrics, confirming the effectiveness

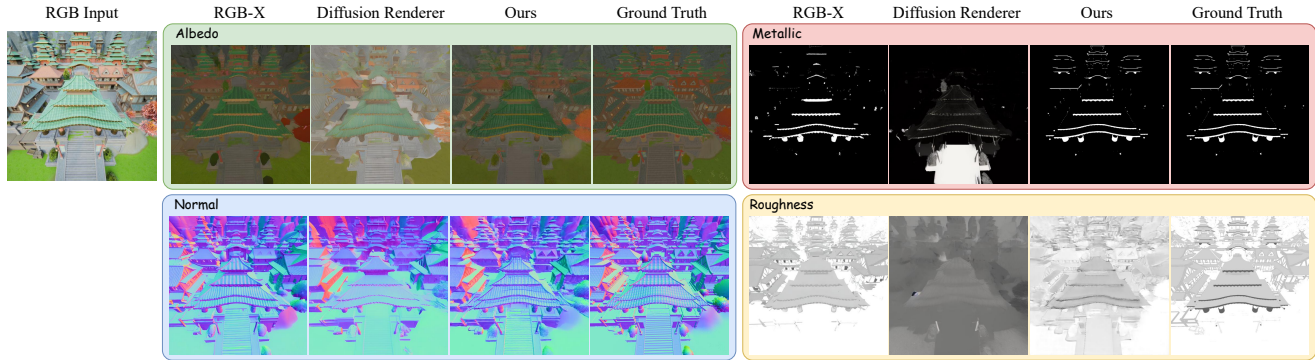


Figure 7. **Qualitative comparison of intrinsic decomposition.** We compare our adapter-based method against RGB-X and DiffusionRenderer. This figure demonstrates the decomposition results for key G-buffer components: **Albedo**, **Normal**, **Metallic**, and **Roughness**.

Method	Paradigm	Params	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Time (s) \downarrow
Path Tracing	Traditional	-	-	-	-	>10
IBL	Traditional	-	16.903	0.882	0.105	-
Deferred	Traditional	-	24.649	0.927	0.097	-
RGB-X [37]	Diffusion	950M	20.984	0.793	0.165	\sim 2.19
DRender [16]	Diffusion	1.7B	23.758	0.863	0.128	\sim 1.40
Ours (w/o key)	Flow	1.4B	24.214	0.874	0.113	\sim 0.19
Ours (w/ key)	Flow	1.7B	26.663	0.883	0.101	\sim 0.24

Table 1. **Quantitative comparison (512x512).** Our method achieves SOTA quality, while being \sim 10x faster than diffusion baselines. Keyframe guidance further improves fidelity with negligible speed impact.

of the proposed guidance mechanism. Qualitative visualizations in Fig. 1 further support these findings. We also report results from Unreal Engine’s native deferred rendering pipeline. While our learned model is not intended as a direct replacement for highly-optimized industry pipelines that operate on explicit high-polygon geometry, the experimental results show that our proposed learning based approach can faithfully reproduce challenging effects such as reflections that are expensive in real-time rendering pipelines. A critical advantage of our approach is its deterministic nature: as illustrated in Fig. 6, our method exhibits zero variance across multiple inference runs, in contrast to the stochastic diffusion baselines, which show significant variance in their outputs. This determinism yields consistent and reproducible output, making our method more reliable for production environments and other applications requiring stable, predictable outputs.

Inference Speed. We also highlight the significant efficiency gains of our method over both iterative baselines and traditional path tracing in Tab. 1. While path tracing times can range from seconds to hours depending on scene complexity, our image-space approach is decoupled from scene geometry, ensuring a consistent inference speed. Compared to diffusion-based methods, our single-step model renders

Method	Albedo		Normal	Material
	PSNR \uparrow	LPIPS \downarrow	Angular ($^\circ$) \downarrow	RMSE \downarrow
RGB-X [37]	26.24	0.116	46.5	0.080
DRender [16]	20.78	0.182	47.6	0.194
Ours	26.93	0.153	16.2	0.084

Table 2. **Quantitative comparison of inverse rendering.** We compare our method with RGB-X and DiffusionRenderer on intrinsics decomposition. The results demonstrate that our parameter-efficient approach achieves competitive performance across key G-buffer attributes.

a 512x512 frame in approximately 0.19 seconds on a single NVIDIA RTX 4090 GPU, making it over 7x faster than DiffusionRenderer and more than 10x faster than RGB-X. Insight: the primary computational overhead stems from the VAE, with G-buffer encoding and final image decoding taking approximately 0.12s and 0.04s respectively, together accounting for nearly 90% of the total inference time.

Inverse Rendering. We also evaluate our inverse adapter’s ability to decompose images into their underlying G-buffers. Note that we are not intended to defeat specialized inverse rendering methods, but to provide a unified framework for forward methods and inverse rendering. As shown in Tab. 2 and Fig. 7, our method demonstrates competitive performance in intrinsics decomposition compared to RGB-X and DiffusionRenderer, highlighting its versatility.

5.2. Ablation Study

We conduct ablation studies to evaluate the effectiveness of the training strategies of bridge matching, the pixel losses, and the keyframe guidance. All the experiments are conducted on lower resolution (256x256) to reduce the computational cost and trained on $4 \times$ RTX 4090s.

Bridge Matching. We compare different training strategies, including flow matching (ODE) and bridge matching (SDE). Following previous fast diffusion works [4, 26], we compare the model trained with a discrete timesteps sam-

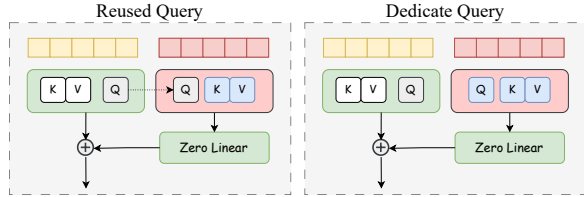


Figure 8. **Two different keyframe adapter designs.** The left one reuses the query from the self-attention layer for efficiency, while the right one uses a dedicated query.

pling scheme for fixed four timesteps $[1.0, 0.75, 0.5, 0.25]$ or in a uniform timesteps. For the SDE-based approach, we follow LBM [17] and set the bridge noise σ to 0.005. As shown in Tab. 3, we empirically find that single-step inference yields superior results for models trained with multi-step schedules, which we attribute to the avoidance of error propagation. Insights: the noise perturbation in the SDE formulation (Eq. 4) encourages the model to generate more diverse effects, thereby increasing its robustness. Ultimately, training with a 4-step SDE schedule and performing inference in a single step provides the best balance of quality and robustness.

Pixel Losses. We evaluate the impact of the pixel-space losses defined in Eq. 7. As detailed in Tab. 4, incorporating an LPIPS loss alongside the latent-space loss significantly improves perceptual quality. The addition of a gradient loss further enhances high-frequency details, such as contact shadows, which are crucial for realistic rendering, leading to the best overall PSNR and SSIM scores.

Keyframe Guidance. We compare several different keyframe adapter designs. We begin with exploring two attention mechanisms with similar ideas to [11] as shown in Fig. 8. The first option is to reuse the query from the self-attention layer, which is the most efficient method of tuning the model for keyframes. Another choice is to use a dedicated query (a full attention layer) for keyframes, which is more flexible but requires more parameters. Moreover, we empirically find that adding a LoRA [9] layer in the feed-forward layer (ffn) in the transformer block can further improve the performance, as it allows the model to adapt to the fused keyframe features. Tab. 5 shows that the dedicated query with LoRA achieves the best performance, outperforming the other two designs. Moreover, Tab. 5, *VACE progressive*, shows that only applying the progressive inference strategy is not sufficient, as the keyframe information is propagated across chunks instead of being directly fed in. Instead, our keyframe adapter provides a more flexible and effective way to incorporate keyframe information, leading to significantly improved rendering quality. The results demonstrate that our method achieves the best performance when combined with the progressive inference strategy. All the experiments are conducted with a keyframe gap

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Uniform SDE (4 steps)	22.192	0.858	0.120
4 timesteps ODE (4 steps)	23.089	0.865	0.110
4 timesteps ODE (1 step)	23.304	0.867	0.108
4 timesteps SDE (4 steps)	23.384	0.865	0.111
4 timesteps SDE (1 step)	23.590	0.868	0.107

Table 3. **Quantitative evaluation of different training strategies.** The method name describes the training and inference setup (e.g., “4 timesteps SDE (1 step)” denotes training with a 4-step SDE schedule but performing inference in a single step).

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
L_{latent} only	21.588	0.840	0.148
$L_{\text{latent}} + L_{\text{lpiPs}}$	23.538	0.867	0.105
$L_{\text{latent}} + L_{\text{lpiPs}} + L_{\text{grad}}$	23.590	0.868	0.107

Table 4. **Quantitative evaluation of the pixel losses.** Experiments demonstrate that the LPIPS loss significantly improves perceptual quality, and the gradient loss further enhances high-frequency details.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
w/o Keyframes	23.590	0.868	0.107
VACE progressive	25.042	0.877	0.102
Reused Query w/o ffn lora	24.520	0.862	0.098
Dedicated Query w/o ffn lora	25.217	0.878	0.097
Dedicated Query w/ ffn lora	25.720	0.881	0.094
Ours + VACE progressive	26.220	0.885	0.090

Table 5. **Quantitative evaluation of the Keyframe Adapter designs.** A dedicated attention query with LoRA in the feedforward layer outperforms the reused one. When combined with the progressive strategies (Sec. 4.4), our full design achieves the highest overall metrics (last row).

of 16 frames. Further analysis on the influence of keyframe gaps is elaborated in the supplementary materials.

6. Conclusion

In this paper, we introduce *RenderFlow*, a novel framework for high-fidelity neural rendering that operates in a single diffusion step. By reformulating the rendering as a conditional flow matching problem, our model learns a direct mapping from G-buffer inputs to the distribution of final images, bypassing the iterative and stochastic sampling procedures of standard diffusion models. When this approach is augmented with a novel sparse keyframe guidance mechanism, *RenderFlow* attains higher visual fidelity than diffusion-based neural rendering baselines while substantially reducing inference time, making it a promising solution for real-time and interactive applications.

References

- [1] Attila T. Áfra. Intel® Open Image Denoise, 2025. <https://www.openimagedenoise.org>. 6
- [2] Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023. 3, 4
- [3] Brent Burley and Walt Disney Animation Studios. Physically-based shading at disney. In *Acm siggraph*, pages 1–7. vol. 2012, 2012. 3
- [4] Clément Chadebec, Onur Tasar, Eyal Benaroché, and Benjamin Aubin. Flash diffusion: Accelerating any conditional diffusion model for few steps image generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 15686–15695, 2025. 7
- [5] Clément Chadebec, Onur Tasar, Sanjeev Sreetharan, and Benjamin Aubin. Lbm: Latent bridge matching for fast image-to-image translation. *arXiv preprint arXiv:2503.07535*, 2025. 2, 5
- [6] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Muller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024. 2
- [7] Kai He, Ruofan Liang, Jacob Munkberg, Jon Hasselgren, Nandita Vijaykumar, Alexander Keller, Sanja Fidler, Igor Gilitschenski, Zan Gojcic, and Zian Wang. Unirelight: Learning joint decomposition and synthesis for video relighting. *arXiv preprint arXiv:2506.15673*, 2025. 2
- [8] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 3
- [9] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022. 8
- [10] Henrik Wann Jensen. *Realistic image synthesis using photon mapping*. Ak Peters Natick, 2001. 1
- [11] Zeyinzi Jiang, Chaojie Mao, Ziyuan Huang, Ao Ma, Yiliang Lv, Yujun Shen, Deli Zhao, and Jingren Zhou. Res-tuning: A flexible and efficient tuning paradigm via unbinding tuner from backbone. *Advances in Neural Information Processing Systems*, 36:42689–42716, 2023. 8
- [12] Zeyinzi Jiang, Zhen Han, Chaojie Mao, Jingfeng Zhang, Yulin Pan, and Yu Liu. Vace: All-in-one video creation and editing. *arXiv preprint arXiv:2503.07598*, 2025. 4, 5
- [13] Haian Jin, Yuan Li, Fujun Luan, Yuanbo Xiangli, Sai Bi, Kai Zhang, Zexiang Xu, Jin Sun, and Noah Snavely. Neural gaffer: Relighting any object via diffusion. *Advances in Neural Information Processing Systems*, 37:141129–141152, 2024. 4
- [14] Brian Karis. Real Shading in Unreal Engine 4. 2013. 3
- [15] Black Forest Labs, Stephen Batifol, Andreas Blattmann, Frederic Boesel, Saksham Consul, Cyril Diagne, Tim Dockhorn, Jack English, Zion English, Patrick Esser, Sumith Kulal, Kyle Lacey, Yam Levi, Cheng Li, Dominik Lorenz, Jonas Müller, Dustin Podell, Robin Rombach, Harry Saini, Axel Sauer, and Luke Smith. Flux.1 kontext: Flow matching for in-context image generation and editing in latent space, 2025. 2
- [16] Ruofan Liang, Zan Gojcic, Huan Ling, Jacob Munkberg, Jon Hasselgren, Zhi-Hao Lin, Jun Gao, Alexander Keller, Nandita Vijaykumar, Sanja Fidler, and Zian Wang. Diffusion-renderer: Neural inverse and forward rendering with video diffusion models. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 2, 4, 6, 7
- [17] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022. 2, 3, 8
- [18] Cheng Ma, Yongming Rao, Yean Cheng, Ce Chen, Jiwen Lu, and Jie Zhou. Structure-preserving super resolution with gradient guidance. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7769–7778, 2020. 5
- [19] Oliver Nalbach, Elena Arabadzhiyska, Dushyant Mehta, H-P Seidel, and Tobias Ritschel. Deep shading: convolutional neural networks for screen space shading. In *Computer graphics forum*, pages 65–78. Wiley Online Library, 2017. 2
- [20] Gaurav Parmar, Taesung Park, Srinivasa Narasimhan, and Jun-Yan Zhu. One-Step Image Translation with Text-to-Image Models, 2024. 2, 5
- [21] Matt Pharr, Wenzel Jakob, and Greg Humphreys. *Physically based rendering: From theory to implementation*. MIT Press, 2023. 1
- [22] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence*, 44(3):1623–1637, 2020. 5
- [23] Erik Reinhard, Michael Stark, Peter Shirley, and James Ferwerda. Photographic tone reproduction for digital images. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 661–670. 2023. 4
- [24] Haocheng Ren, Yuchi Huo, Yifan Peng, Hongtao Sheng, Hongxiang Huang, Weidong Xue, Jingzhen Lan, Rui Wang, and Hujun Bao. Lightformer: Light-oriented global neural rendering in dynamic scene. *ACM Transactions on Graphics*, 43(4):1–14, 2024. 2
- [25] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 1
- [26] Axel Sauer, Frederic Boesel, Tim Dockhorn, Andreas Blattmann, Patrick Esser, and Robin Rombach. Fast high-resolution image synthesis with latent adversarial diffusion distillation. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–11, 2024. 7
- [27] Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. Adversarial diffusion distillation. In *European Conference on Computer Vision*, pages 87–103. Springer, 2024. 2

- [28] Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger bridge matching. *Advances in Neural Information Processing Systems*, 36: 62183–62223, 2023. [3](#)
- [29] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. [4](#)
- [30] Manu Mathew Thomas and Angus G Forbes. Deep illumination: Approximating dynamic global illumination with generative adversarial network. *arXiv preprint arXiv:1710.09834*, 2017. [2](#)
- [31] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, Jianyuan Zeng, Jiayu Wang, Jingfeng Zhang, Jingren Zhou, Jinkai Wang, Jixuan Chen, Kai Zhu, Kang Zhao, Keyu Yan, Lianghua Huang, Mengyang Feng, Ningyi Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang, Tianxing Wang, Tianyi Gui, Tingyu Weng, Tong Shen, Wei Lin, Wei Wang, Wei Wang, Wenmeng Zhou, Wenteng Wang, Wenting Shen, Wenyuan Yu, Xianzhong Shi, Xiaoming Huang, Xin Xu, Yan Kou, Yangyu Lv, Yifei Li, Yijing Liu, Yiming Wang, Yingya Zhang, Yitong Huang, Yong Li, You Wu, Yu Liu, Yulin Pan, Yun Zheng, Yuntao Hong, Yupeng Shi, Yutong Feng, Zeyinzi Jiang, Zhen Han, Zhi-Fan Wu, and Ziyu Liu. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025. [2](#), [4](#)
- [32] Jay Zhangjie Wu, Yuxuan Zhang, Haithem Turki, Xuanchi Ren, Jun Gao, Mike Zheng Shou, Sanja Fidler, Zan Gojcic, and Huan Ling. Difix3d+: Improving 3d reconstructions with single-step diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 26024–26035, 2025. [2](#)
- [33] Rongyuan Wu, Lingchen Sun, Zhiyuan Ma, and Lei Zhang. One-step effective diffusion network for real-world image super-resolution. *Advances in Neural Information Processing Systems*, 37:92529–92553, 2024. [2](#)
- [34] Jingjing Xu, Xu Sun, Zhiyuan Zhang, Guangxiang Zhao, and Junyang Lin. Understanding and improving layer normalization. *Advances in neural information processing systems*, 32, 2019. [4](#)
- [35] Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Fredo Durand, William T Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6613–6623, 2024. [2](#)
- [36] Chong Zeng, Yue Dong, Pieter Peers, Hongzhi Wu, and Xin Tong. Renderformer: Transformer-based neural rendering of triangle meshes with global illumination. In *ACM SIGGRAPH 2025 Conference Papers*, 2025. [2](#)
- [37] Zheng Zeng, Valentin Deschaintre, Iliyan Georgiev, Yannick Hold-Geoffroy, Yiwei Hu, Fujun Luan, Ling-Qi Yan, and Miloš Hašan. Rgbx: Image decomposition and synthesis using material- and lighting-aware diffusion models. In *ACM SIGGRAPH 2024 Conference Papers*, New York, NY, USA, 2024. Association for Computing Machinery. [2](#), [5](#), [6](#), [7](#)
- [38] Howard Zhang, Yuval Alaluf, Sizhuo Ma, Achuta Kadambi, Jian Wang, and Kfir Aberman. Instantrestore: Single-step personalized face restoration with shared-image attention, 2024. [2](#)
- [39] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. [5](#), [6](#)