

# Supplementary Materials of BSDF importance sampling using a diffusion model

## 1 ODE EQUATION DERIVATION

We will validate ODE Equation  $F(x_t, t) = \mathbb{E}[x_1 - x_0 | x_t, t]$  purely from probabilistic perspective. We denote the corresponding density of  $x_t$  as  $p_t$ .

For clarity and coherence in our exposition, we will maintain predefined notations for marginal probability densities, such as  $p(x_0) = p_0$ . However, to simplify our presentation in the following sections, we will not assign unique notations to each joint probability distribution. Instead, we use intuitive notation; for instance,  $p(x_0, x_1)$  denotes the joint probability density function of  $x_0$  and  $x_1$ .

Starting from the analytical form of expectation, we express the conditional expectation as follows:

$$\mathbb{E}[x_1 - x_0 | x_t, t] = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} (x_1 - x_0) p(x_0, x_1 | x_t, t) dx_0 dx_1.$$

Applying Bayes' theorem, we obtain:

$$p(x_0, x_1 | x_t, t) = \frac{p(x_t | x_0, x_1, t) p(x_0) p(x_1)}{p(x_t | t)}.$$

Substituting into the expectation, we derive:

$$\mathbb{E} = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} (x_1 - x_0) \frac{p(x_t | x_0, x_1, t) p(x_0) p(x_1)}{p(x_t | t)} dx_0 dx_1.$$

The function  $p(x_t | x_0, x_1, t)$  behaves essentially as a Dirac delta function, activated when  $x_t = (1-t)x_0 + tx_1$ . This allows us to express  $x_0$  as  $x_0 = \frac{x_t - tx_1}{1-t}$ .

Refining further, we arrive at:

$$\mathbb{E} = \frac{1}{p_t(x_t | t)} \int_{-\infty}^{+\infty} \frac{1}{1-t} \frac{x_1 - x_t}{1-t} p_0\left(\frac{x_t - tx_1}{1-t}\right) p_1(x_1) dx_1, \quad (1)$$

where  $\frac{1}{1-t}$  is the Jacobian of the change of variables. We define the transformations:

$$x_0^t = (1-t)x_0 \sim \frac{1}{1-t} p_0\left(\frac{x_0}{1-t}\right) = p_0^t,$$

where  $p_0^t$  approaches a Dirac delta function as  $t \rightarrow 1$ . Similarly, we define

$$x_1^t = tx_1 \sim \frac{1}{t} p_1\left(\frac{x_1}{t}\right) = p_1^t,$$

where  $p_1^t$  converges to a Dirac delta function as  $t \rightarrow 0$ . This leads to  $x_t = x_0^t + x_1^t$ , with the corresponding density  $p_t$  being the convolution of  $p_0^t$  and  $p_1^t$ .

The analytical expression for  $p_t(x_t | t)$  is:

$$\begin{aligned} p_t(x_t | t) &= \int_{-\infty}^{+\infty} p_0^t(x_t - x_1) p_1^t(x_1) dx_1 \\ &= \int_{-\infty}^{+\infty} \frac{1}{1-t} p_0\left(\frac{x_t - x_1}{1-t}\right) \frac{1}{t} p_1\left(\frac{x_1}{t}\right) dx_1. \end{aligned}$$

Subsequently, we define the cumulative distribution function (CDF) of  $p_t$  as  $c_t$ , given by:

$$c_t(z | t) = \int_{-\infty}^z \int_{-\infty}^{+\infty} \frac{1}{1-t} p_0\left(\frac{x_t - x_1}{1-t}\right) \frac{1}{t} p_1\left(\frac{x_1}{t}\right) dx_1 dx_t.$$

Author's address:




Fig. 1. Straight paths can reduce the number of sampling steps.

Knowing that  $\frac{dx_t}{dt} = \frac{dx_t}{dc_t} \cdot \frac{dc_t}{dt}$  and recognizing that  $\frac{dc_t}{dx_t}$  is the density function  $p_t(x_t | t)$ , we focus on deriving  $\frac{dc_t}{dt}$ . Setting  $k = \frac{x_1}{t}$ , we transform  $c_t(z | t)$  as follows:

$$c_t(z | t) = \int_{-\infty}^z \int_{-\infty}^{+\infty} \frac{1}{1-t} p_0\left(\frac{x_t - tk}{1-t}\right) p_1(k) dk dx_t$$

Then we do another change of variables  $m = \frac{x_t - tk}{1-t}$ :

$$c_t(z | t) = \int_{-\infty}^{\frac{z-tk}{1-t}} \int_{-\infty}^{+\infty} p_0(m) p_1(k) dk dm.$$

The derivative  $\frac{d \frac{z-tk}{1-t}}{dt} = \frac{k-z}{(1-t)^2}$ , applying Leibniz's rule, results in:

$$\frac{d}{dt} c_t(z | t) = \int_{-\infty}^{+\infty} \frac{k-z}{(1-t)^2} p_0\left(\frac{z-tk}{1-t}\right) p_1(k) dk.$$

Replacing  $k$  with  $x_1$  and  $z$  with  $x_t$ , we conclude:

$$\frac{dx_t}{dt} = \frac{dx_t}{dc_t} \cdot \frac{dc_t}{dt} = \frac{1}{p_t(x_t | t)} \int_{-\infty}^{+\infty} \frac{x_1 - x_t}{(1-t)^2} p_0\left(\frac{x_t - tx_1}{1-t}\right) p_1(x_1) dx_1.$$

This equation exactly matches the analytical form of expectation as given in Equation (1). Through the proof outlined above, it is evident that theoretically, the ODE learns by modeling a time-continuous set of distributions  $p_t$ , which is the convolution of two other time-continuous distributions  $p_0^t$  and  $p_1^t$ . This fundamentally aligns with the original concept of the diffusion model, except that it replaces noise with an arbitrary known distribution and eliminates the need for complex Evidence Lower Bound (ELBO) calculations, directly utilizing samples to compute expectations as the values that the network needs to predict.

Consequently, the loss function for our diffusion model employed is as follows:

$$\min_{\theta} \mathbb{E}_{t, x_t} \left[ \left\| D_{\theta}(x_t, t) - \mathbb{E}_{(x_0, x_1) | (x_t, t)} [x_1 - x_0] \right\|^2 \right] \quad (2)$$

Where  $D_{\theta}$  is a neural network, mostly represented by small MLP in our case.

Furthermore, since the distribution  $p_t$  results from a convolution, the presence of the integral operation ensures that even if  $p_0$  or  $p_1$  are not continuous on  $\mathbb{R}^d$ , the resulting learned distribution remains continuous on  $\mathbb{R}^d$ . This inherent characteristic limits the diffusion model's ability to capture discontinuities. To provide a




Fig. 2. Evolution of  $x_t$  over time  $t$  during the learning process of mapping between two uniform distributions using the diffusion model. The red areas represent histograms generated from actual samples, while the blue dots denote the pdf values for each sample. These pdf values are computed by multiplying the base samples' pdf values by the Jacobian derived from the ODE.




Fig. 3. Some results for 1D distributions that satisfy the conditions presented. When the distributions are continuous, despite their complexity and multimodality, the diffusion model can still yield accurate outcomes.

clearer illustration, consider a scenario using the diffusion model to learn a mapping from one uniform distribution to another of the same type.

Based on our theoretical derivation,  $p_t$  evolves as the convolution of two distributions over time. For uniform distributions, visualizing the shape of their convolution is straightforward because it essentially involves the convolution of two uniform distributions, with only the defined interval changing. Figure 2 confirms two key conclusions previously discussed: (1) The learned distribution  $p_t$  is a convolution of two distributions, and (2) such type of diffusion models is inherently incapable of accurately capturing discontinuities. At  $t = 0.5$ , the result is clearly the convolution of two identical uniform distributions, which increasingly approximates the target uniform distribution as time progresses. However, since the learning process involves convolution, the pdf values must start and end at zero. This requirement results in extremely sharp changes at the boundaries, when  $t$  comes to 1. Aside from cases where the target distribution itself is inherently discontinuous, most scenarios we encounter involve distributions that exhibit discontinuities at their boundaries. This occurs because convolution is defined over the domain  $\mathbb{R}$ . For distributions like the uniform distribution or the BRDFs that we plan to explore, these are theoretically continuous within their intrinsic domains and manifest discontinuities only at the boundaries. Figure 3 demonstrates the effectiveness of the diffusion model when all conditions are met.

## 2 BIJECTION

The deterministic ODE version of the diffusion model offers attractive and essential properties. For the ODE model, in theory it




Fig. 4. The one-to-one mapping learned by the diffusion model.




Fig. 5. Crossings are impossible in the ODE trajectories; if they were to occur, it would imply that at some  $x_t$ , the ODE points in two different directions.

supports learning a bijective mapping between a simple base distribution and a complex target distribution, which is exactly what normalizing flows achieve. The bijective property is crucial for our use of the diffusion model for sampling. Without bijectivity, it means the mapping we learn is not invertible, making our calculation of the jacobian invalid. The bijective nature of the mappings used in diffusion models has been confirmed by prior studies [??]. We provide an intuitive explanation. Given that trajectories are derived using an ODE, and assuming a well-defined ODE where  $f$  is a continuous function, it is impossible for these trajectories to intersect. If intersections were possible, there would exist a set  $(x_t, t)$  for which the value of  $f(x_t, t)$  would not be unique, contradicting our assumptions. Furthermore, this non-intersection property allows us to infer that the mapping learned by the diffusion model resembles the mapping used in **inverse transform sampling**, and it will be exactly inverse transform sampling once you set the base distribution to  $U(0, 1)$ . Specifically, if  $x'_0$  evolves through the ODE to  $x'_1$ , then it holds that  $c_0(x'_0) = c_1(x'_1) = u$ , where  $c_0$  and  $c_1$  CDFs of  $p_0$  and  $p_1$  respectively.

## 3 TIME STATISTICS

Our model training was conducted on a single RTX 4090 GPU, consisting of three distinct stages: pretraining (10,000 epochs), diffusion (40,000 epochs), and rectification (10,000 epochs). The majority of computational time was consumed by the diffusion and rectification stages. Notably, while the rectification stage demonstrated rapid convergence, typically requiring only a few thousand epochs, its time consumption remained substantial due to the necessity of performing hundreds of diffusion model sampling steps per epoch. The sample generation process, executed on a single-core AMD Ryzen 7 5800X CPU, required approximately 0.5 hours to produce 4.9 million samples.

Domain	Pretrain	Diffusion	Rectify
Disk	16 min	41 min	1 h 22 min
Spherical	18 min	59 min	1 h 35 min

Table 1. Training times for different Domains across stages.

#### 4 PDF EVALUATION

In the sampling process of diffusion model, each discrete step in the ODE can be conceptualized as an invertible transform. This perspective allows us to compute the Jacobian determinant for each step, which is then multiplied by the PDF  $p_0$  of the base samples  $x_0$ . As the ODE values are predicted using a neural network, the network’s derivatives with respect to the input  $x_t$  are easily obtainable and then can be used to compute the Jacobian determinant. The final PDF for  $x_1$  conditioned on  $\omega_i$  is given by  $p(\omega_o|\omega_i) = J_1 * p_0(x_0|\omega_i)$ , where  $J_1$  represents the accumulated Jacobian computed through the forward ODE trajectory during the sampling process, formulated as follows:

$$J_{t+\frac{1}{N}} = J_t \left( 1 + \frac{1}{N} * \left| \det \left( \frac{dD_\theta(\hat{x}_t, t, \omega_i)}{d\hat{x}_t} \right) \right| \right)$$

for  $t$  ranging over  $\{0, \dots, N\}/N$ , with the initial value  $J_0$  set to one.

To facilitate Multiple Importance Sampling, it is also necessary to compute the PDF for the reverse trajectory, tracing from the BSDF  $p(\omega_o | \omega_i)$  back to the base  $p_0$ . The diffusion model supports efficient the reverse process, which takes the exact same steps as the forward one. We reverse the trajectory to obtain the base samples  $x_0$ , then we calculate the final PDF for  $\omega_o$  conditioned on  $\omega_i$  is given by  $p(\omega_o|\omega_i) = J_0 * p_0(x_0|\omega_i)$ . The the accumulated Jacobian  $J_0$  for reverse case:

$$\hat{x}_{1-t-\frac{1}{N}} = \hat{x}_{1-t} - \frac{D_\theta(\hat{x}_{1-t}, 1-t, \omega_i)}{N}$$

$$J_{1-t-\frac{1}{N}} = J_{1-t} * \left( 1 - \frac{1}{N} * \left| \det \left( \frac{dD_\theta(\hat{x}_{1-t}, 1-t, \omega_i)}{d\hat{x}_{1-t}} \right) \right| \right)$$

where  $t \in \{0, \dots, N\}/N$ , with  $\hat{x}_1 = \omega_o$  are inputs from the BSDF,  $J_1$  initial value is one.

#### 5 PRETRAIN NETWORK

Given different  $\omega_o$ , the corresponding BRDF distributions can vary significantly in shape and variance, as Fig 6 (a) shows. This preliminary step proves beneficial in reducing the size of the diffusion model’s network and the number of sampling steps required, particularly when dealing with energy-concentrated case, mainly specular materials and around grazing angle where Fig. 6 (b) provides one example Fig. 6 (b) provides one example. Subsequently, we use this pretrained distribution as the base distribution for the diffusion model, with  $\omega_o$  serving as the target distribution, conditional on  $\omega_i$ .

In our practical tests, using an extremely complex distribution does not necessarily enhance the expressiveness of the diffusion model. We experimented with a Gaussian mixture with 4 lobes as the base distribution and found it to be less effective than using a uniform distribution, particularly in anisotropic specular cases. We

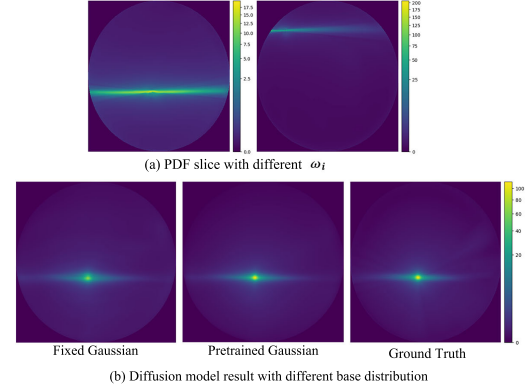


Fig. 6. (a) PDF slices of a single material from the RGL dataset, demonstrating variations in variance and shape with different incoming directions. (b) Comparison of the diffusion model when trained using a unit Gaussian and a pretrained Gaussian as the base distribution respectively, with the same sampling steps and network size.

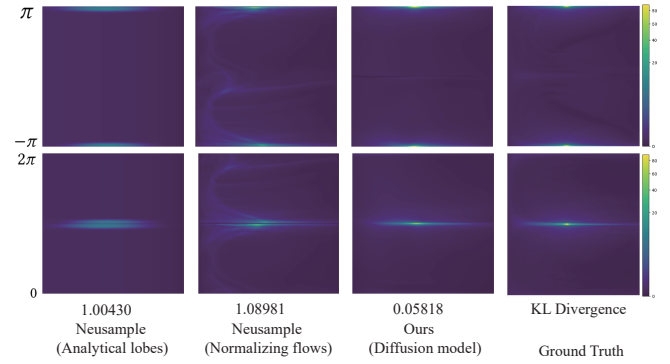


Fig. 7. PDF slices of the METAL-PAPER-COPPER material are examined with a fixed incoming direction and a specified range for  $\theta$ . The upper  $\phi$  values range from  $[-\pi, \pi]$ , and the lower  $\phi$  values from  $[0, 2\pi]$ . We compare NeuSample’s mixture lobe and Normalizing Flows methods results with our method. By accurately handling the periodicity on the  $\phi$  channel, our approach can effectively learn scenarios where there are high PDF values at the boundaries.

believe this is because Gaussian lobes can become highly concentrated in highlights, and due to anisotropy, they may concentrate in incorrect positions, making training more difficult. Therefore, we opted for a small network that roughly captures the shape of the distribution while keeping the base distribution relatively diffuse. A Gaussian distribution proved to be a very reasonable choice.

#### 6 PDF SLICES

We present a comprehensive comparison of PDF slices across a broader range of materials, contrasting these with the NeuSample [Xu et al. 2023] on disk domain. Additionally, we compare these results with our earlier diffusion model prior to implementing ReFlow, to demonstrate the enhanced accuracy achieved with ReFlow.

*Periodicity in the Spherical Domain.* Considering periodicity in the spherical domain is crucial. Figure 7 illustrates a special case where

highlights at the boundary fail to be learned properly without considering periodicity due to discontinuities. However, by accounting for periodicity, the distribution is perceived as continuous, allowing the highlights at the boundary to be handled correctly and producing reasonable results.

And since previous method does not support spherical domain sampling, we only compare our method results with ground truth. The disk domain PDF slices are from Figure 19 to Figure 24, while spherical domain PDF slices are from Figure 25 to Figure 29.

## 7 RENDERING RESULTS

In this section, we present comparison between our model and NeuSample [Xu et al. 2023] and NBRDF [Sztrajman et al. 2021] in terms of image rendering, focusing on equal SPP (samples per pixel) to illustrate the expressiveness of each model. Under point light illumination, such as Figure 9, 11, indirect lighting significantly contributes to rendering noise, overshadowing BSDF sampling effects. Conversely, in ambient lighting conditions, such as 8, 10, BSDF sampling predominantly influences rendering outcomes. This motivates our focus on soft ambient lighting scenarios for experimental evaluation. The NBRDF model, which utilizes the isotropic Blinn-Phong reflectance model, demonstrates satisfactory performance

for materials with low anisotropy. However, it exhibits significant limitations when applied to highly anisotropic materials, resulting in considerable noise artifacts in the generated results.

For diffuse materials all sampling methods effectively capture the distribution, resulting in almost no variance.

For conductor materials that are not extremely specular the NeuSample method begins to exhibit weaker performance around the grazing angle. In contrast, our model remains stable.


For extremely specular materials the NeuSample method fails completely. Although our method shows a higher occurrence of fireflies on the disk, the spherical domain demonstrates excellent stability and accurately captures the distribution at grazing angles.

We also observe that, overall, diffusion models within the spherical domain exhibit a slower convergence speed. This is particularly evident for materials with strong anisotropic properties.


## REFERENCES

- Alejandro Sztrajman, Gilles Rainer, Tobias Ritschel, and Tim Weyrich. 2021. Neural BRDF representation and importance sampling. *Computer Graphics Forum* 40, 6 (2021), 332–346.
- Bing Xu, Liwen Wu, Milos Hasan, Fujun Luan, Iliyan Georgiev, Zexiang Xu, and Ravi Ramamoorthi. 2023. NeuSample: Importance Sampling for Neural Materials. In *ACM SIGGRAPH 2023 Conference Proceedings*. 1–10.


Ground Truth




NBRDF




Neusample Analytical



Neusample Normflows



Ours Disk



Ours Spherical




Fig. 8. Rendering results at 128 spp using environmental illumination on real-world captured materials.

Ground Truth



NBRDF




Neusample Analytical



Neusample Normflows



Ours Disk



Ours Spherical



Fig. 9. Rendering results at 128 spp using point illumination on real-world captured materials.




Fig. 10. Rendering results at 128 spp using environmental illumination on real-world captured materials. , Vol. 1, No. 1, Article . Publication date: September 2024.

Ground Truth



NBRDF



Neusample Analytical



Neusample Normflows



Ours Disk



Ours Spherical







Fig. 12. Rendering results at 128 spp using environmental and point illumination on synthesized BSDFs. Vol. 1, No. 1, Article . Publication date: September 2024.

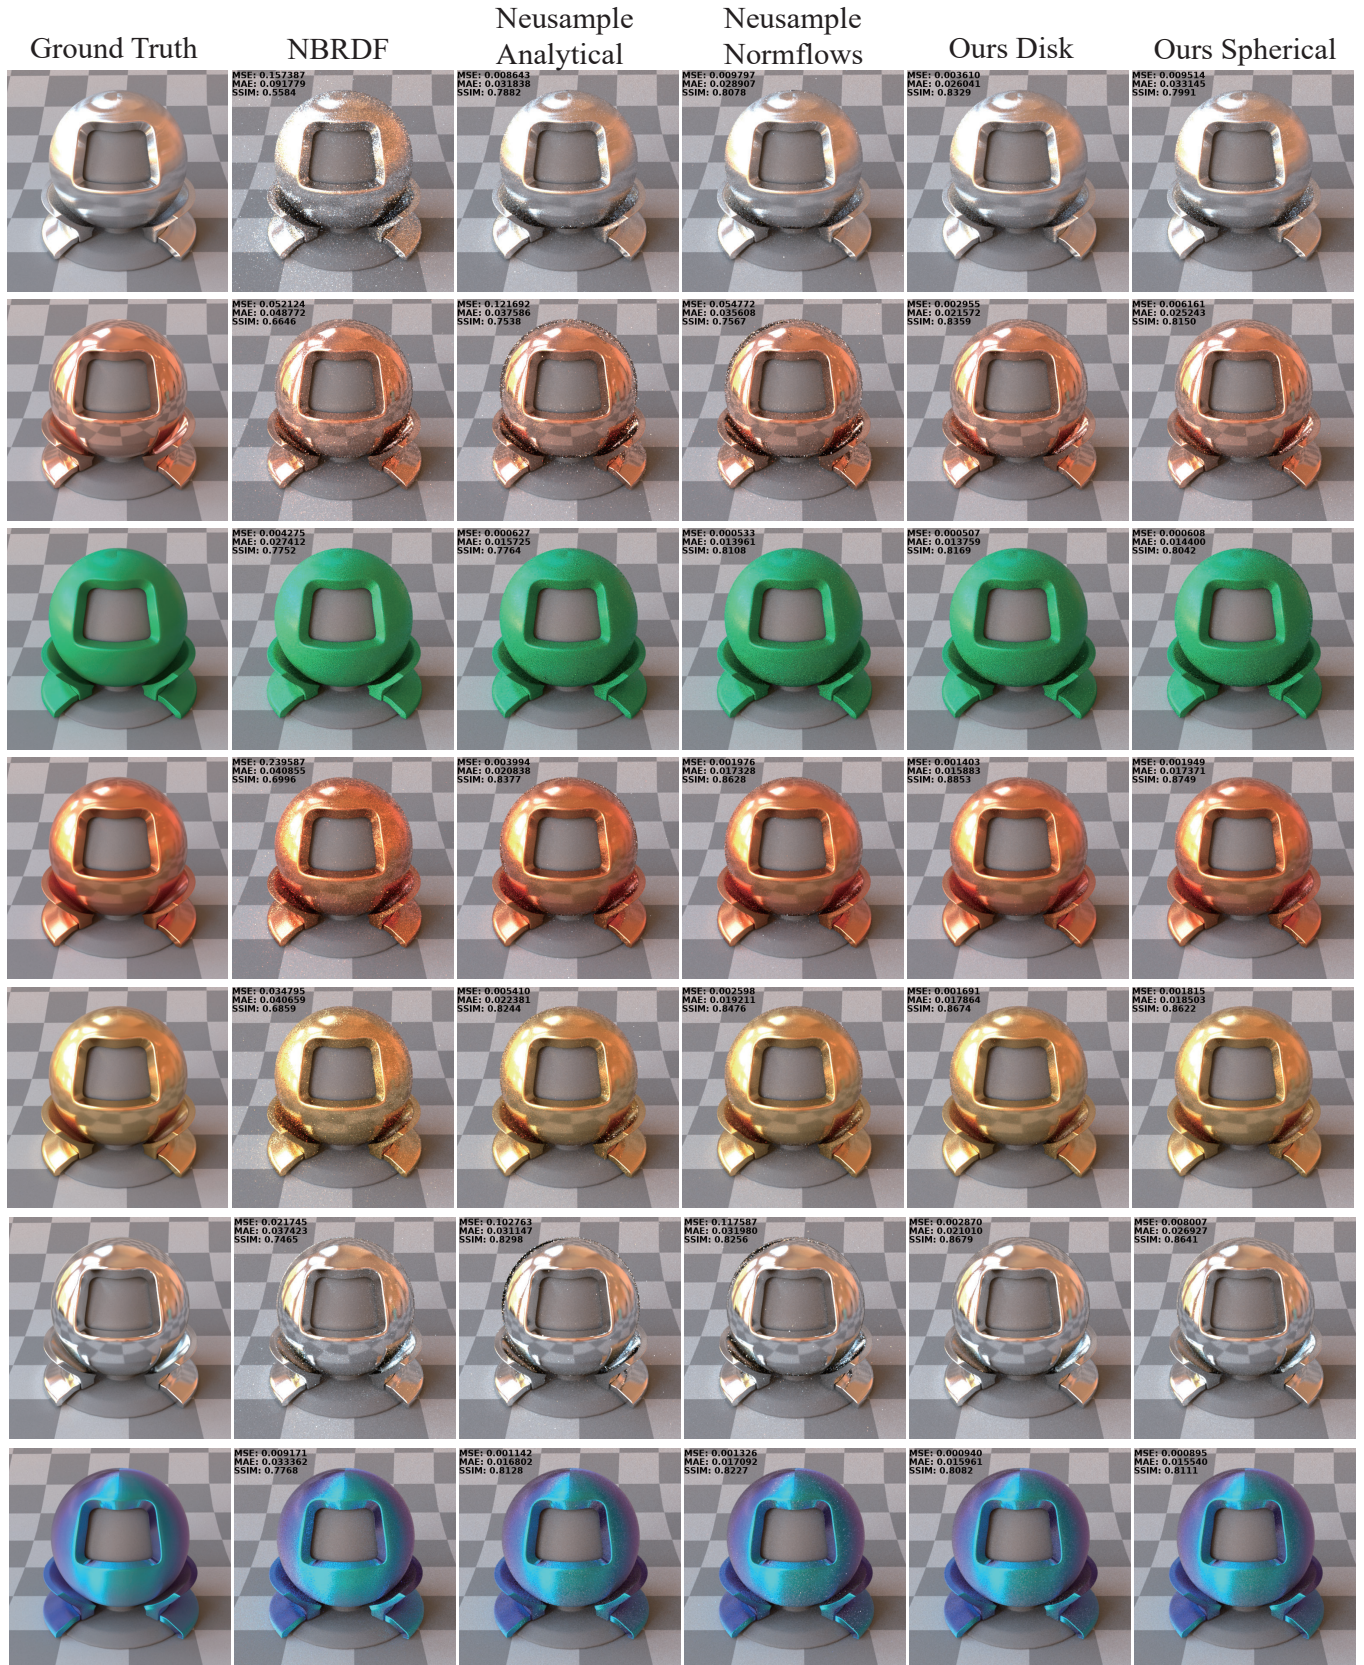


Fig. 13. Rendering results at 128 spp on real-world captured materials.



Fig. 14. Rendering results at 128 spp on real-world captured materials.



Fig. 15. Rendering results at 128 spp on real-world captured materials.




Fig. 16. Rendering results at 128 spp on real-world captured materials.



Fig. 17. Rendering results at 128 spp on real-world captured materials.



Fig. 18. Rendering results at 128 spp on synthesized BSDF. Vol. 1, No. 1, Article . Publication date: September 2024.

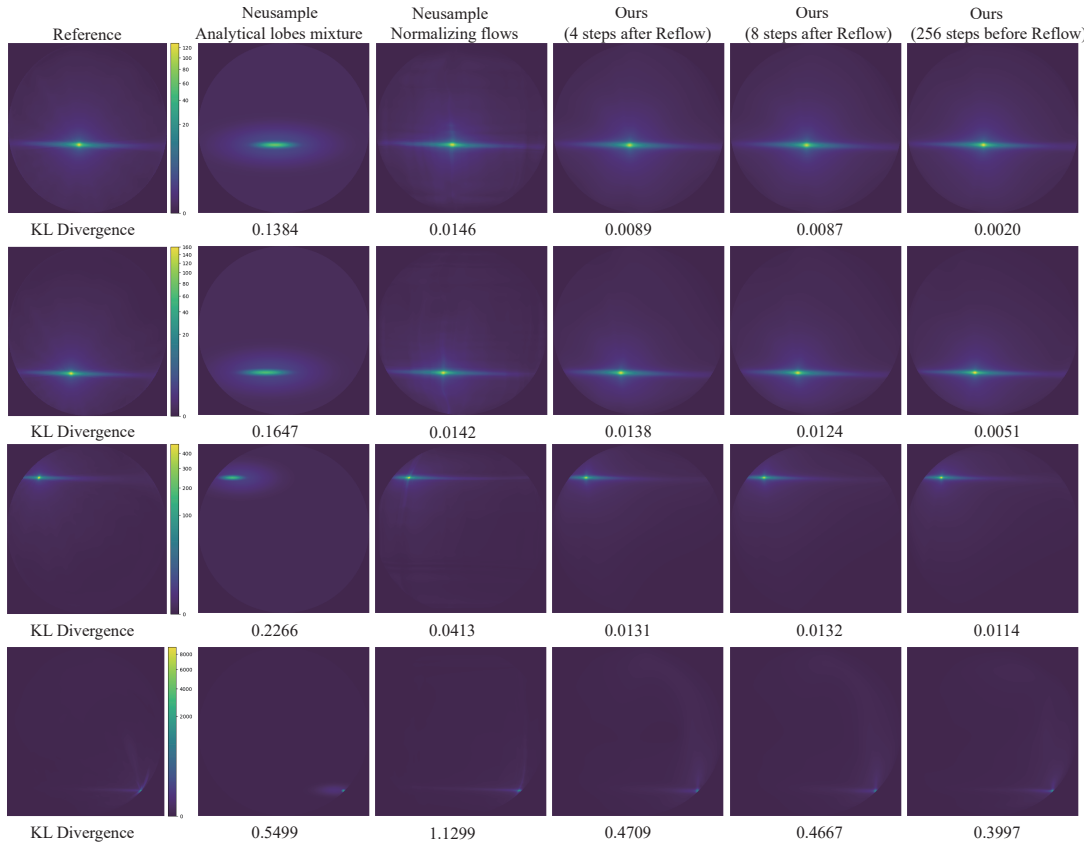


Fig. 19. Disk domain pdf slices on METAL-PAPER-COPPER.

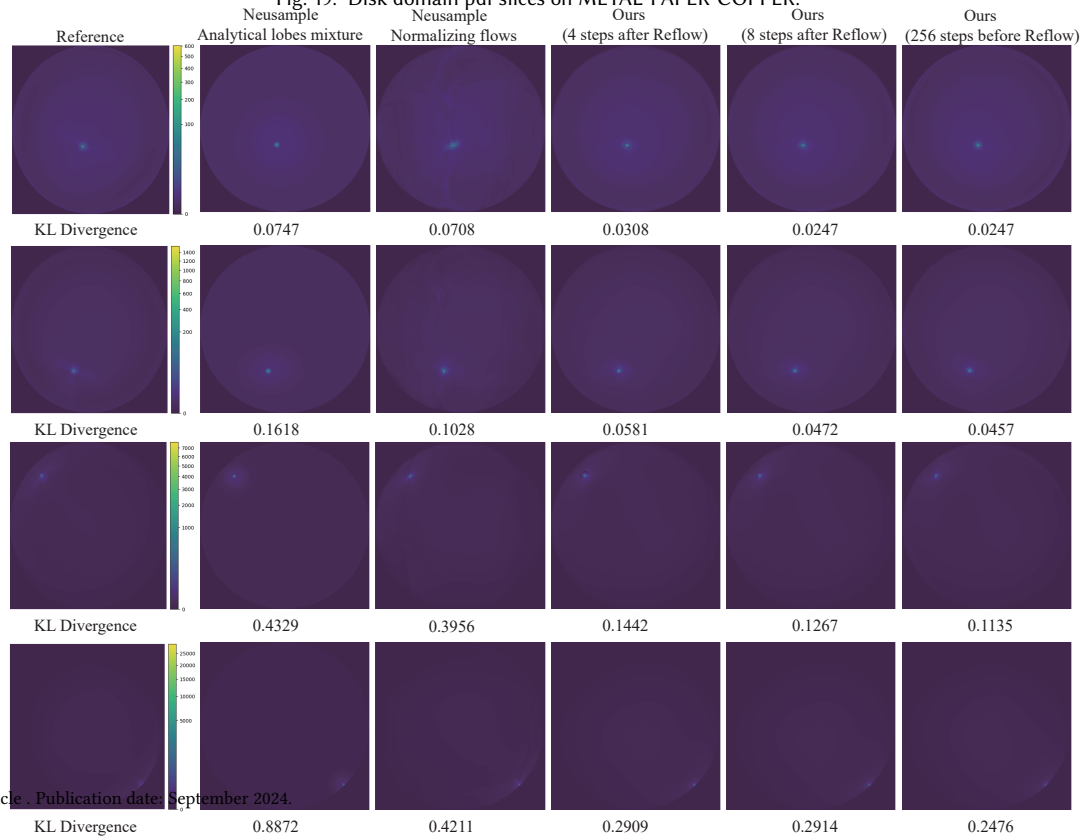


Fig. 20. Disk domain pdf slices on GREEN-MALACHITE.






Fig. 21. Disk domain pdf slices on MORPHO-MELENAUS.




Fig. 22. Disk domain pdf slices on SLIKE-BLUE.




Fig. 23. Disk domain pdf slices on IBIZA-SUNSET.

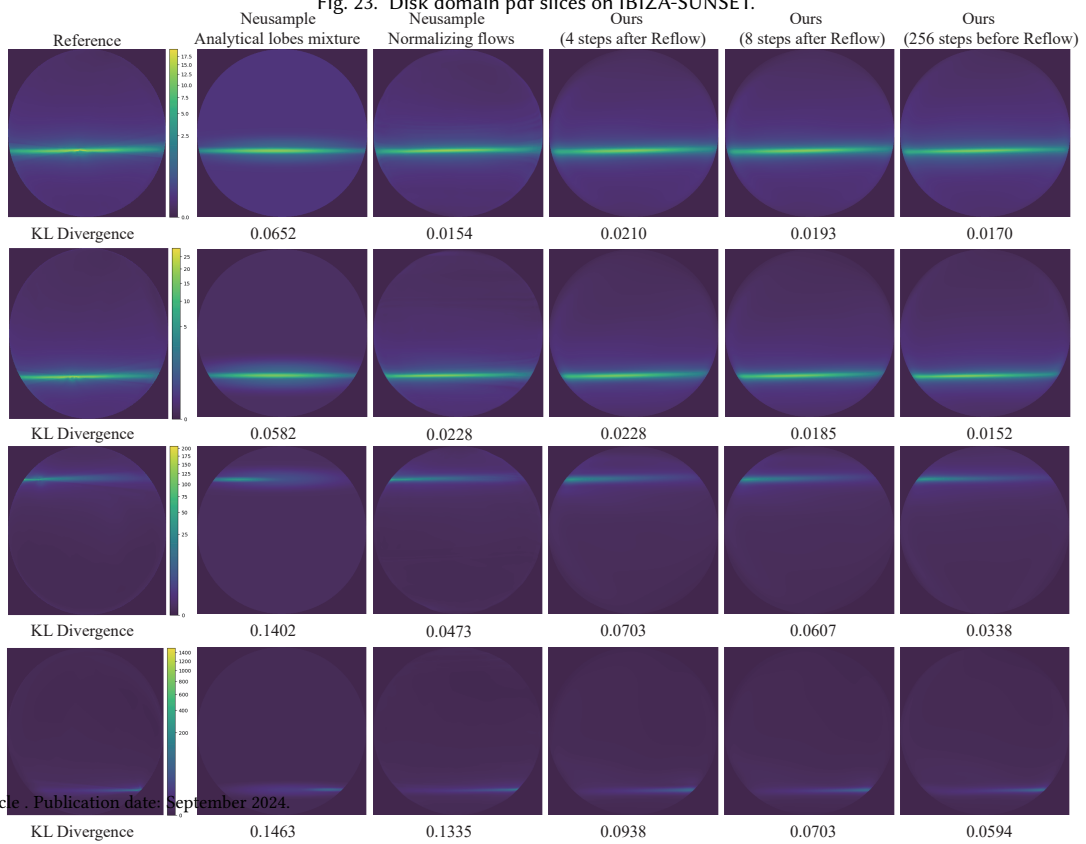


Fig. 24. Disk domain pdf slices on BRUSHED-STEEL-SATIN-PINK.

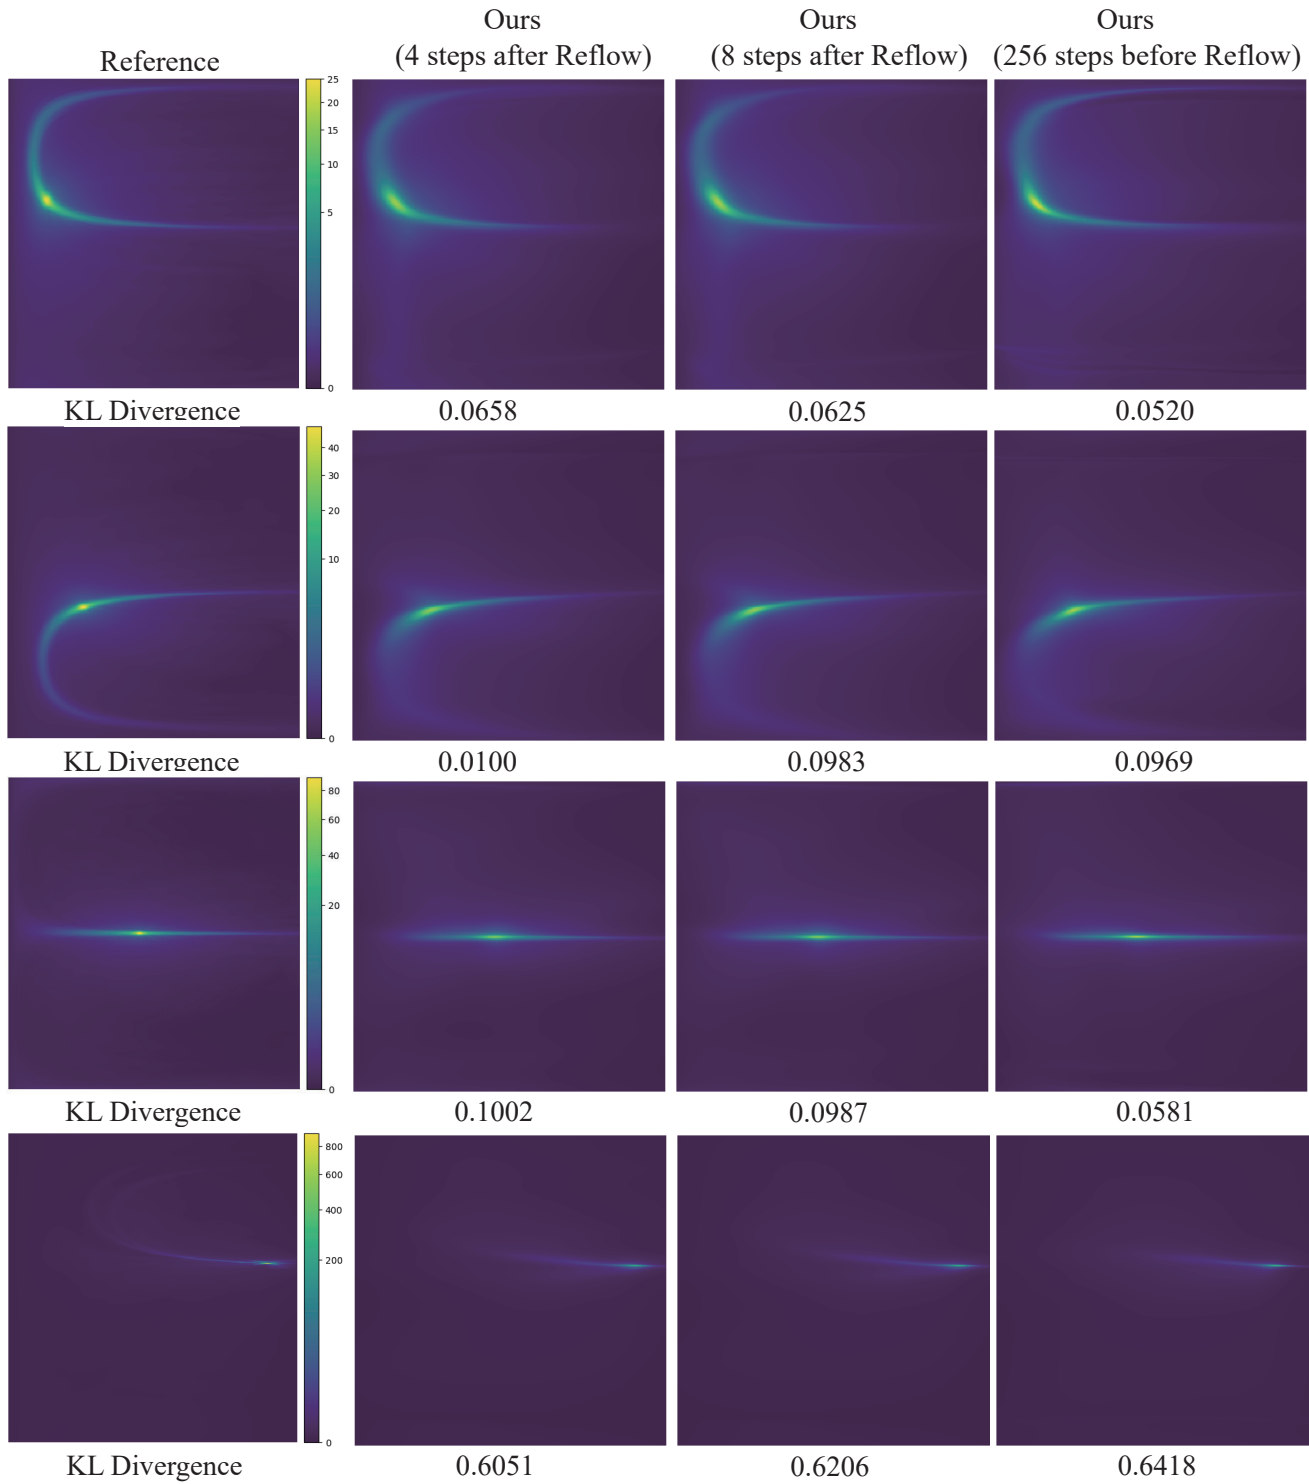


Fig. 25. Spherical domain pdf slices on METAL-PAPER-COPPER.




Fig. 26. Spherical domain pdf slices on GREEN-MALACHITE.




Fig. 27. Spherical domain pdf slices on MORPHO-MELENAUS.




Fig. 28. Spherical domain pdf slices on SLIKE-BLUE.



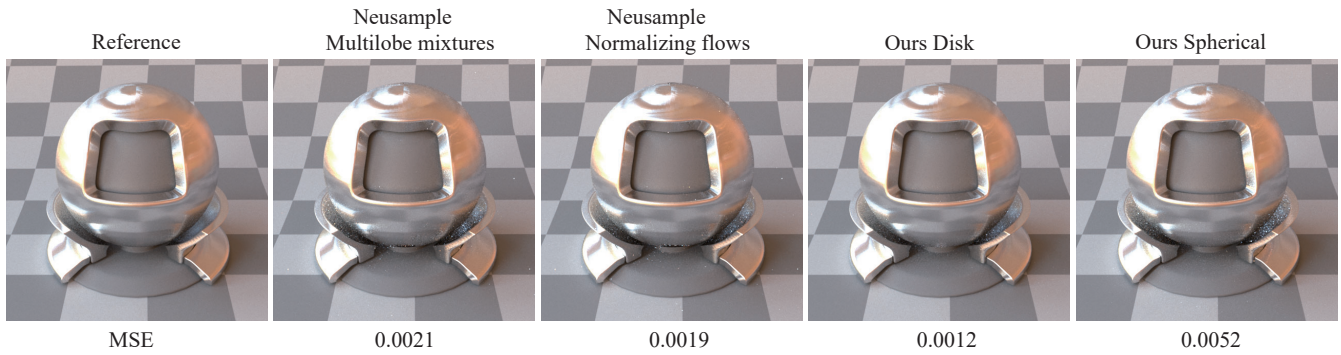


Fig. 29. Spherical domain pdf slices on BRUSHED-STEEL-SATIN-PINK.




(a) Only BRDF sampling




(b) Multiple importance sampling

Fig. 30. Results at 516 SPP Using BRDF-Only Sampling and MIS on ANISO-BRUSHED-ALUMINIUM.




(a) Only BRDF sampling




(b) Multiple importance sampling

Fig. 31. Results at 516 SPP Using BRDF-Only Sampling and MIS on COPPER-PAPER-METAL.






(a) Only BRDF sampling




(b) Multiple importance sampling

Fig. 32. Results at 516 SPP Using BRDF-Only Sampling and MIS on COPPER-SHEET.




(a) Only BRDF sampling




(b) Multiple importance sampling

Fig. 33. Results at 516 SPP Using BRDF-Only Sampling and MIS on GOLD-PAPER-METAL.




(a) Only BRDF sampling

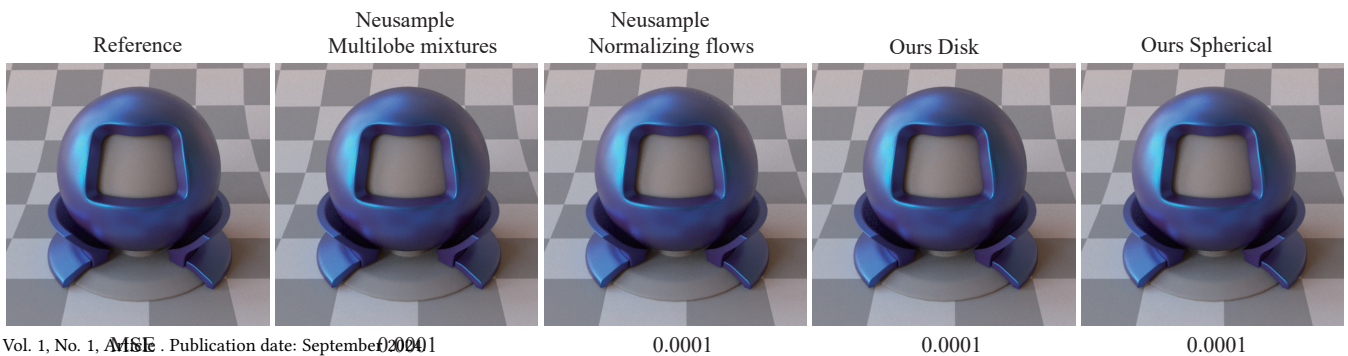


(b) Multiple importance sampling

Fig. 34. Results at 516 SPP Using BRDF-Only Sampling and MIS on GREEN-MALACHITE.

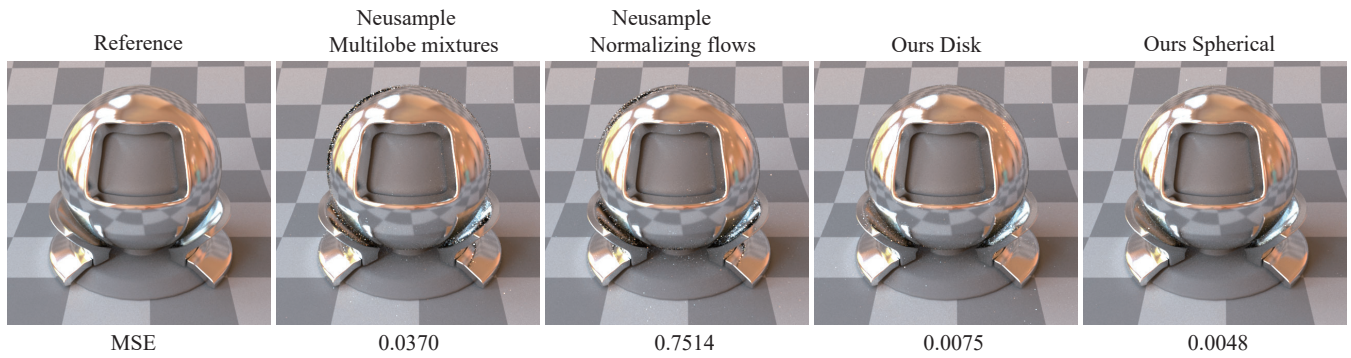


(a) Only BRDF sampling

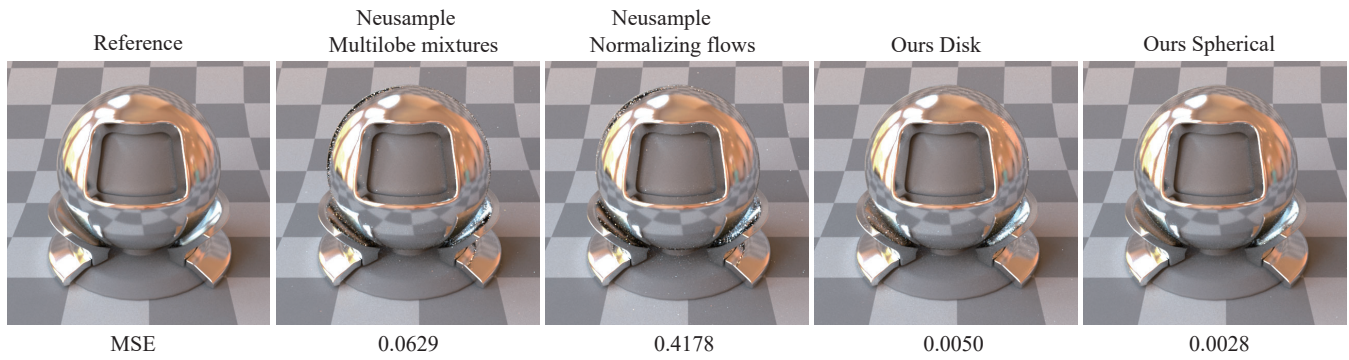


(b) Multiple importance sampling

Fig. 35. Results at 516 SPP Using BRDF-Only Sampling and MIS on SILK-BLUE.

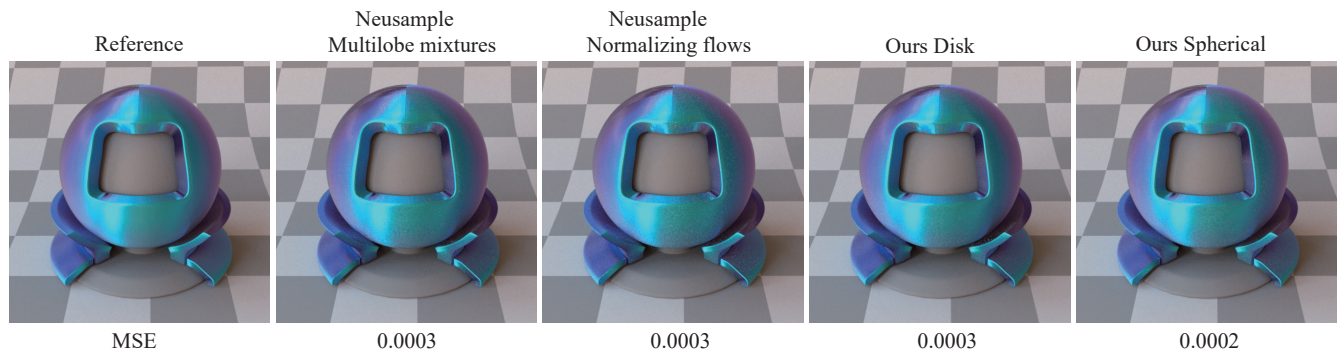


(a) Only BRDF sampling

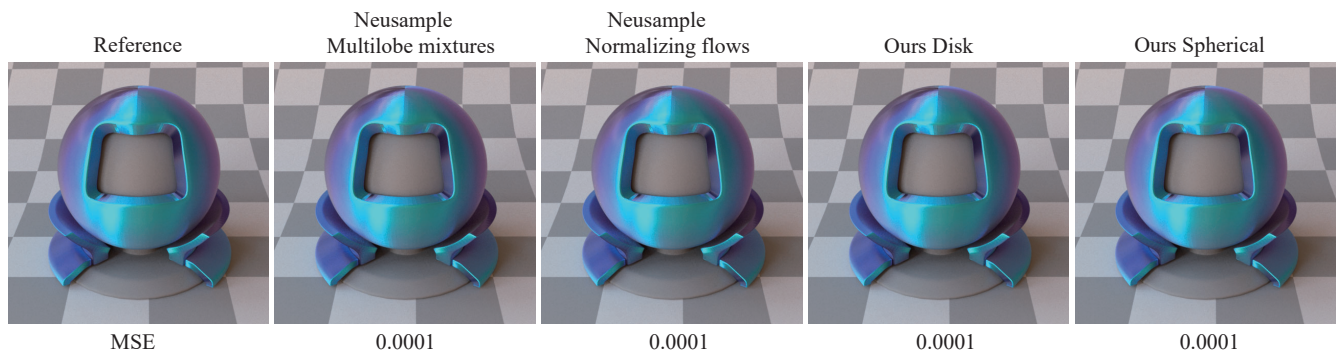


(b) Multiple importance sampling

Fig. 36. Results at 516 SPP Using BRDF-Only Sampling and MIS on ANISO-MIRRO.

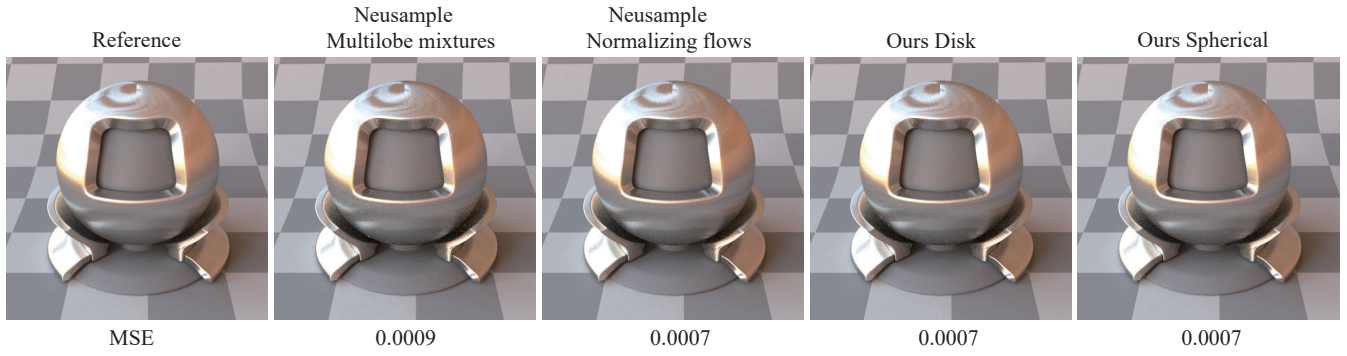


(a) Only BRDF sampling

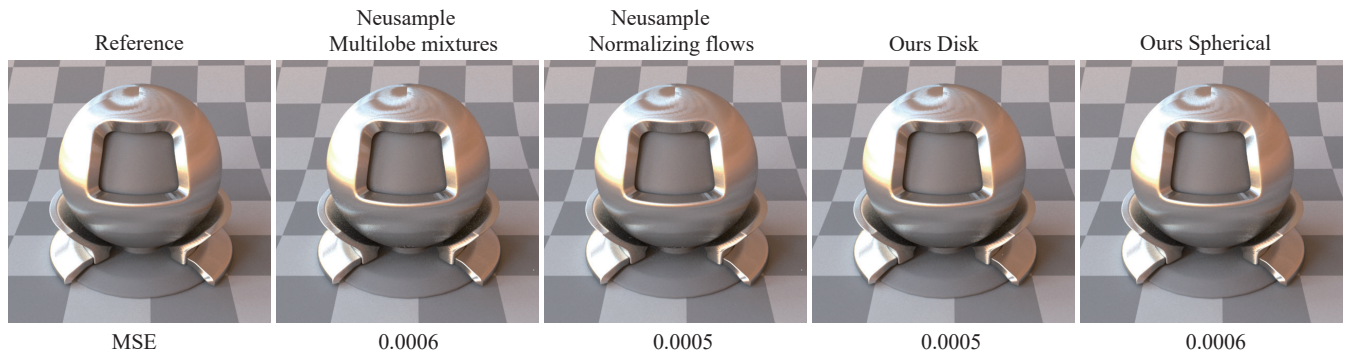


(b) Multiple importance sampling

Fig. 37. Results at 516 SPP Using BRDF-Only Sampling and MIS on ANISO-MORPHO-MELENAUS.



(a) Only BRDF sampling



(b) Multiple importance sampling

Fig. 38. Results at 516 SPP Using BRDF-Only Sampling and MIS on BRUSHED-STEEL-SATIN-PINK.