Multiple Scattering using Machine Learning

Feng Xie
Stanford University

Anton Kaplanyan
Facebook Reality Labs

Warren Hunt
Facebook Reality Labs

Pat Hanrahan
Stanford University

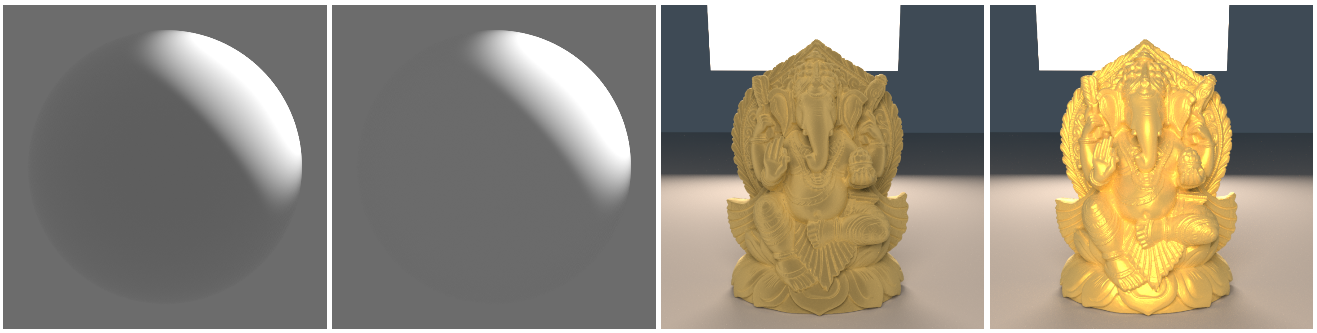

(a) Single scattering (b) Gaussian Multi-scattering (c) Single scattering (d) RealNVP Multi-scattering

Figure 1: Images of a white sphere and a gold figurine rendered with and without multiple scattering at roughness $\alpha = 0.7$ using the Beckmann Distribution. Both multiple scattering images are able to preserve energy. Additionally the gold figurine rendered with multiple scattering displays color saturation effects.

CCS CONCEPTS

• **Computing methodologies → Reflectance modeling**; **Rendering**;

KEYWORDS

Appearance Modeling, Microfacet Distributions, Multiple Scattering

ACM Reference Format:
Feng Xie, Anton Kaplanyan, Warren Hunt, and Pat Hanrahan. 2019. Multiple Scattering using Machine Learning. In *Proceedings of SIGGRAPH '19 Talks.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3306307.3328144

1 INTRODUCTION

Microfacet-based reflection models are widely used in visual effects applications ranging from computer games to animation and feature film rendering. However, the standard microfacet BRDF supports single scattering only. Light that is scattered more than once is not accounted for, which can lead to significant energy loss. As physically based rendering becoming more prevalent in production, the lack of energy preservation has become problematic. This has lead to several recent works on multiple scattering. Heitz et al. [2016] presented a volumetric approach to model multiple scattering accurately but its stochastic evaluation increases variance. Xie and Hanrahan [2018] presented an analytical multiple scattering model that is efficient but has a singularity in the direction of mirror reflection.

In this talk, we present two novel approaches for modeling multiple reflections from specular micro-surfaces. Figure 1 illustrates the visual impact of including microfacet inter-reflections. In the left two images, a simple white sphere is lit by a grey dome and an area light. Compared to single scattering only, the sphere rendered with Gaussian fitted multiple scattering is brighter due to improved energy preservation. In the right two images, a gold figurine rendered with Real NVP based multiple scatterig not only preserves energy but also demonstates color saturation effects due to multiple scattering.

2 SIMULATION OF MULTIPLE SCATTERING

To simulate single and multiple reflection on a rough surface, we first create a Gaussian heightfield surface for a roughness value α. The normals of a Gaussian heightfield surface follow the Beckmann distribution. For each incident angle θ_i, we shoot n rays against the heightfield. [1] We trace each ray until it exits the hemisphere around the heightfield, then record the final exit direction ω_o and number of reflections. This is henceforth referred to as the MCRT simulation.

In our machine learning based approaches, we aim to fit or model the slope distribution of the half vector ω_h computed by the MCRT simulation. The slope for $\omega_h = (x_h, y_h, z_h)$ is defined as $\hat{n}_h = \left(\frac{x_h}{z_h}, \frac{y_h}{z_h}\right)$. Then we compute the distribution of ω_o using the chain rule: $p(\omega_o) = p(\hat{n}_h) \left|\det\left(\frac{\partial \hat{n}_h}{\partial \omega_o}\right)\right|$ where $\left|\det\left(\frac{\partial \hat{n}_h}{\partial \omega_o}\right)\right| = \frac{\omega_i \cdot \omega_o + 1}{(z_i + z_o)^3}$.

[1] n is in the order of millions.

SIGGRAPH '19 Talks, July 28 - August 01, 2019, Los Angeles, CA, USA

ACM ISBN 978-1-4503-6317-4/19/07.
https://doi.org/10.1145/3306307.3328144

3 GAUSSIAN FITTING OF MULTIPLE SCATTERING

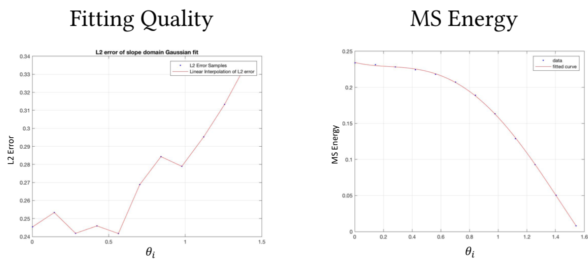

Figure 2: L2-norm of Gaussian fitting and multiple scattering energy ratio for $\alpha = 0.7$.

Since the slope distribution for a given θ_i is uni-modal, our first approach is to fit the distribution of multiple scattering $\hat{h}_n$ using a parametrized $2D$ Gaussian $N(x; \mu(\theta_i), \sigma(\theta_i))$.

For each uniformly sampled θ_i in $(0, \frac{\pi}{2})$, we use the EM method to fit a $2D$ Gaussian for the MCRT simulated $\hat{n}_h$ distribution. Then we use low degree polynormials to fit μ and σ, the coefficients of the $2D$ Gaussians computed at discrete samples of θ_i.

The effective L2 error of our Gaussian fitting is consistently lower than 5% due to the inverse relationship between the fitting error and the energy ratio of multiple scattering. An example is shown in Figure 2 for $\alpha = 0.7$. Since the distribution of $\hat{n}_h$ changes smoothly with θ_i, low degree polynomial fitting of μ and σ works well. Detailed analysis of the $2D$ Gaussian and polynomial fitting are reported in the supplementary material.

We validate the visual quality of this approach by creating a GaussianBrdf object inside PBRT. The multiple scattering image in Figure 1b demonstrates this Gaussian fitted model is able to preseve the reflected energy better than single scattering; moreover, it enhances the shape of the reflected highlights without introducing any artifacts.

4 BRDF MODELING USING REAL NVP

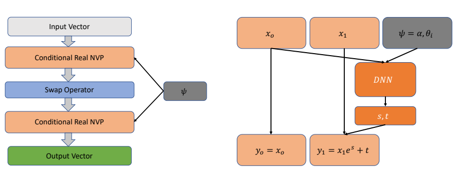

(a) (b)

Figure 3: Design of a Conditional Real NVP Chain and a Conditional Real NVP Layer

The second approach we present for modeling the multiple scattering BRDF comes from the field of variational inferencing. Dinh et al. [2016] presented Real NVP, a set of invertible and learnable transformations between any data distribution p_x and a simple latent distribution p_z. They apply this technique to learn the transformation between a high dimentional Gaussian and a multi-modal distribution of images. We are the first to apply this technique to model the multiple scattering of rough surfaces.

4.1 Model Design and Training

The scale and translation factors (s, t) of the standard Real NVP model depend only on the latent input vector. However, for our application, the distribution of $\hat{n}_h$ also depends on θ_i and α; but not vice versa. To express this dependency, we extend the definition for (s, t) to include conditional dependency on an external parameter $\psi = (\theta_i, \alpha)$ that is independent of the latent input. ψ can be augmented to include other spatially varying parameters like Fresnel or anisotropy.

Figure 3a shows the design of our transform network *RealNVPChain*. The input is a multivariate diagonal Gaussian $N(x; \alpha)$. The chain itself is composed of n layers of conditional bijector shown in Figure 3b interlaced with a swap bijector $(y_0, y_1) = (x_1, x_0)$ to ensure all components are covered by the transform network. For training, we use a loss function that maximizes the likelihood of the observed simulation data:

$$loss = -\sum \text{target\_dist.log\_prob}(simulation\_data) \tag{1}$$

4.2 Results

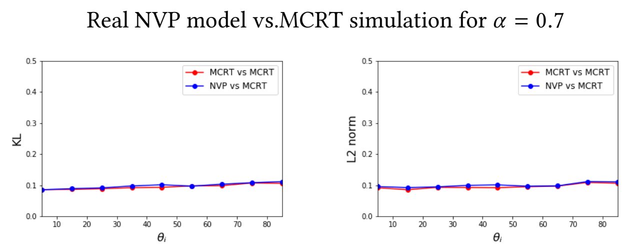

Figure 4: Kullback-Leibler Divergence and the L2 norm between Real NVP distribution and MCRT simulation.

Trained conditional Real NVP is quite effective at modeling the distribution of $\hat{n}_h$. Our experiments use a chain with 4 layers where each layer uses a DNN of 30 nodes. Figure 4 shows the KL divergence and L2 norm between our Real NVP model and the MCRT simulation are within 5% of the corresponding variance measures between two consequtive MCRT simulations. Details of the Real NVP model and results are presented in the supplementary material.

To validate the visual quality of our model, we integrate the trained model as a MLBrdf object in PBRT using TensorFlow serving. The multiple scattering gold figurine image in Figure 1d is rendered using the MLBrdf. Compared to the standard single scattering image, this image preserves energy and demonstrates color saturation effects unique in rough metals. Compared to the parametrized Gaussian fitting, the Real NVP model is more accurate and general. However, the parametrized Gaussian model is easily added to existing rendering systems while the Real NVP model requires system level changes to support efficient DNN inferencing.

REFERENCES

A. P. Dempster, N. M. Laird, and D. B. Rubin. 1977. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* 39, 1 (1977), 1–38. http://www.jstor.org/stable/2984875

Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. 2016. Density estimation using Real NVP. *CoRR* abs/1605.08803 (2016). arXiv:1605.08803 http://arxiv.org/abs/1605.08803

Eric Heitz, Johannes Hanika, Eugene d'Eon, and Carsten Dachsbacher. 2016. Multiple-scattering Microfacet BSDFs with the Smith Model. *ACM Trans. Graph.* 35, 4, Article 58 (July 2016), 14 pages. https://doi.org/10.1145/2897824.2925943

Feng Xie and Pat Hanrahan. 2018. Multiple scattering from distributions of specular v-grooves. *ACM Trans. Graph.* 37, 6 (2018), 276:1–276:14. https://doi.org/10.1145/3272127.3275078