TensoFlow: Tensorial Flow-based Sampler for Inverse Rendering

Supplementary Material

7. More results

We further provide a quantitative analysis of the sample variance in the rendering equation's evaluation. Specifically, the per-pixel variance is defined as:

$$\text{Var} = \frac{1}{N(N-1)} \sum_{i=0}^{N} \frac{(I_i - \overline{I})^2}{q_i}, \tag{25}$$

where I denotes the integrand value of the rendering equation for a sampled incident direction, $\overline{I}$ represents the average integrand value across all samples, and q_i is the probability density function (PDF) value for each sample. For quantitative evaluation, we report the mean variance across all pixels of the image for the specular term in Tab. 4. Here, "Pre-defined" refers to the use of a pre-defined importance sampler (GGX distribution [30]) for the rendering equation. The results clearly demonstrate that our learnable importance sampler achieves significantly lower variance compared to the predefined sampler.

Furthermore, Tab. 5 shows the impact of each component of *TensoFlow* on variance reduction, highlighting the effectiveness of each module in improving the evaluation of the rendering equation.

In Fig. 7, we provide additional visualizations of material decomposition, normal maps, and per-pixel variance on the TensoSDF dataset [20].

Table 4. Quantitative evaluation of variance ↓ (in Units of $1e^{-5}$) in rendering equation evaluation using pre-defined variance and our learnable importance sampler.

	Pre-defined	Ours
Rover	4.389	**1.281**
Dragon	3.382	**1.217**
Motor	11.957	**3.340**
Helmet	2.357	**0.691**
Robot	3.200	**1.922**
Compressor	13.184	**4.413**
Average	6.411	**2.144**

Table 5. Ablation studies on variance ↓ (in Units of $1e^{-5}$) across various components of *TensoFlow*.

	Pre-defined	Ours
w/o half	4.955	**1.773**
w/o tensorial	3.486	**2.679**
w/o reflected	4.301	**1.506**
$N_s = 32$	21.371	**2.780**
$N_s = 64$	9.044	**1.996**
Full	6.411	**2.144**

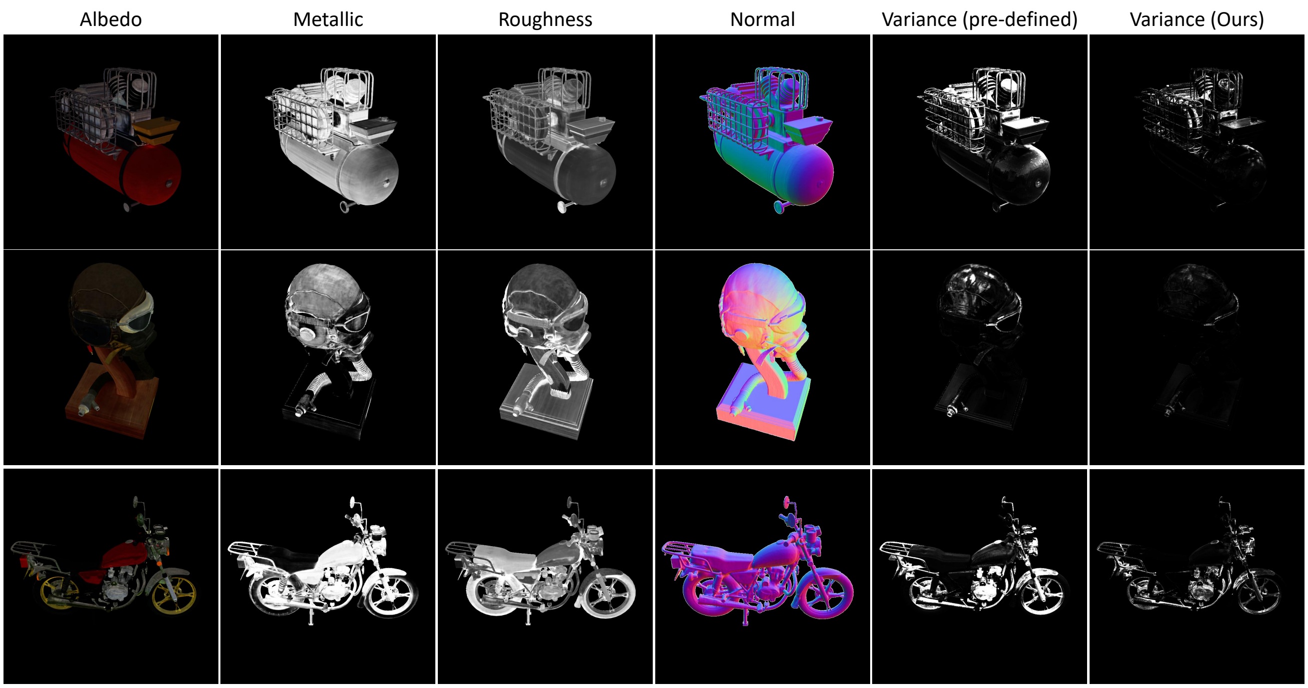

Figure 7. Visualization of material decomposition, normal map, and per-pixel variance in rendering equation evaluation, comparing results with a pre-defined sampler and our proposed tensorial normalizing flow. The per-pixel variance results are scaled by 1000 for clearer visualization.