Rendering of Layered Materials using Statistical Analysis: Supplemental Material

submission id: 153

Contents

1 Smooth Coating

We validated our layered implementation using the Symmetric model against a Ground Truth that stochastically evaluates the layered model. Our first test is a coated sphere lit by two disk lights. We set the layered material to have two layers, a first smooth dielectric layer with varying index of refraction and a rough conducting base with $\boldsymbol{\eta} = 1 + [1, 0.5, 0.5]\,\mathbf{i}$. We display the two rendering side-by-side.

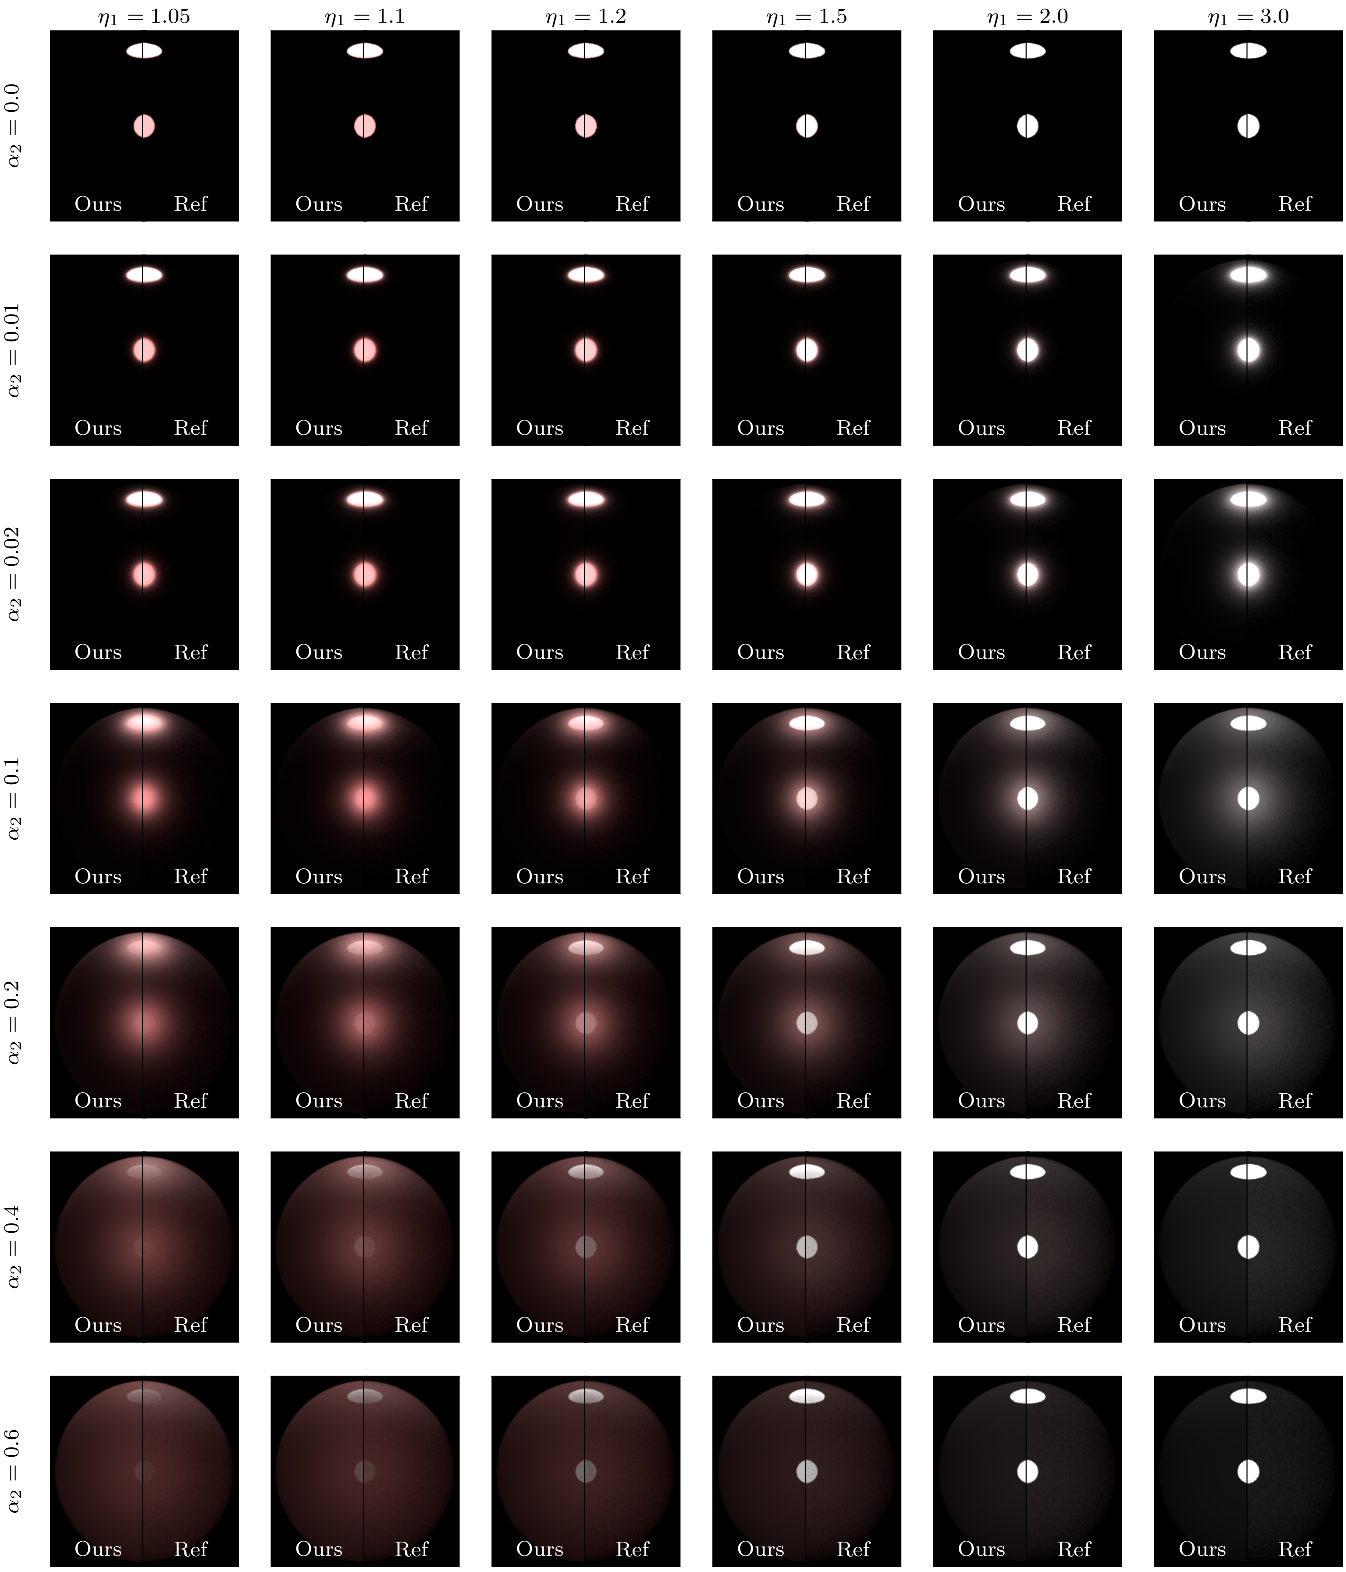

2 Rough Coating

Our second test is a frosted sphere lit by two disk lights. We set the layered material to have two layers, a first rough dielectric layer with varying index of refraction and a smooth conducting base with $\boldsymbol{\eta} = 1 + [1, 0.5, 0.5]\,\mathbf{i}$. We display the two rendering side-by-side.

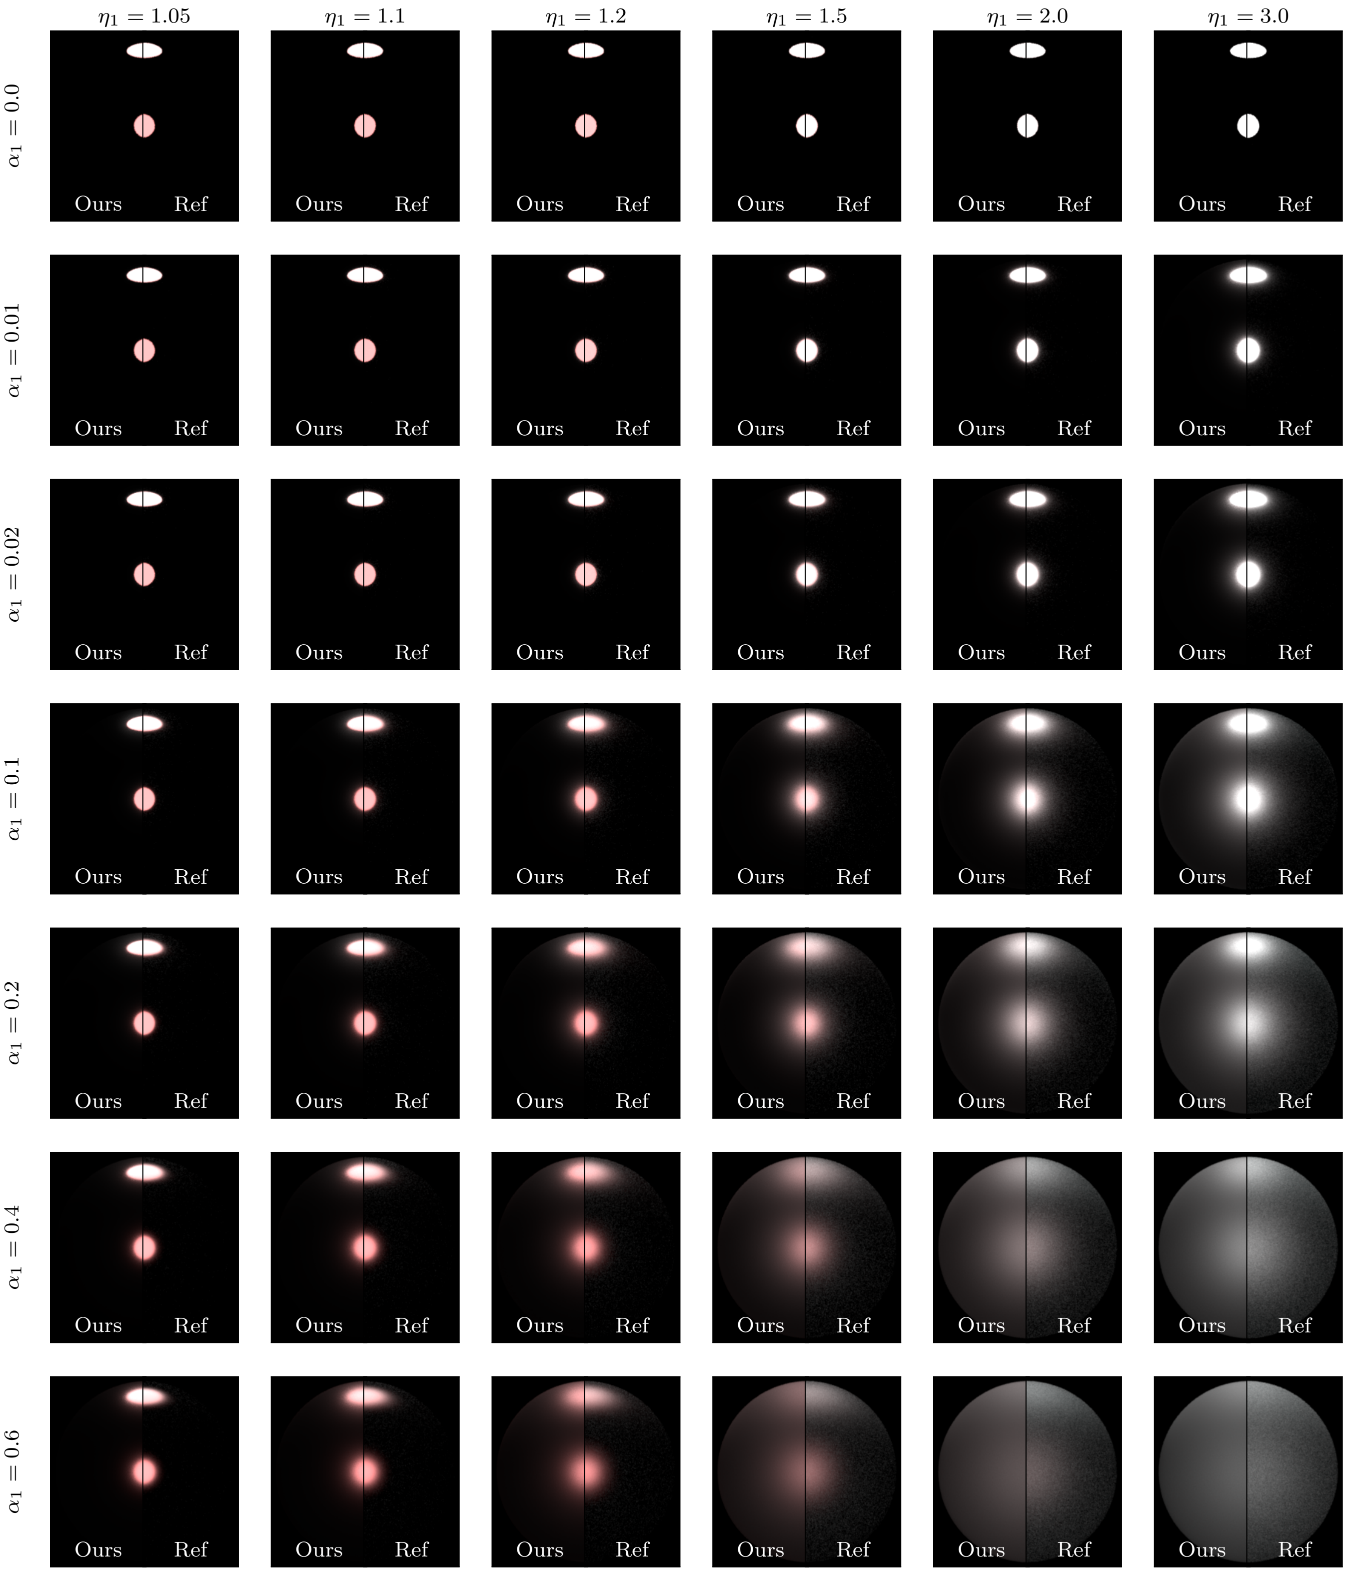

3 Faked Refraction

We validated our faked rough refraction using a simple unit test. We set a tilted plane in front of a perspective camera lit by an environment map emitter. We varied the real index of refraction η and the roughness of the surface α. We display side-by-side the result of our layered BSDF model and of mitsuba `roughdieletric` plugin. Aside from the noise level due to difference in sampling strategies, we found the difference to be minimal.

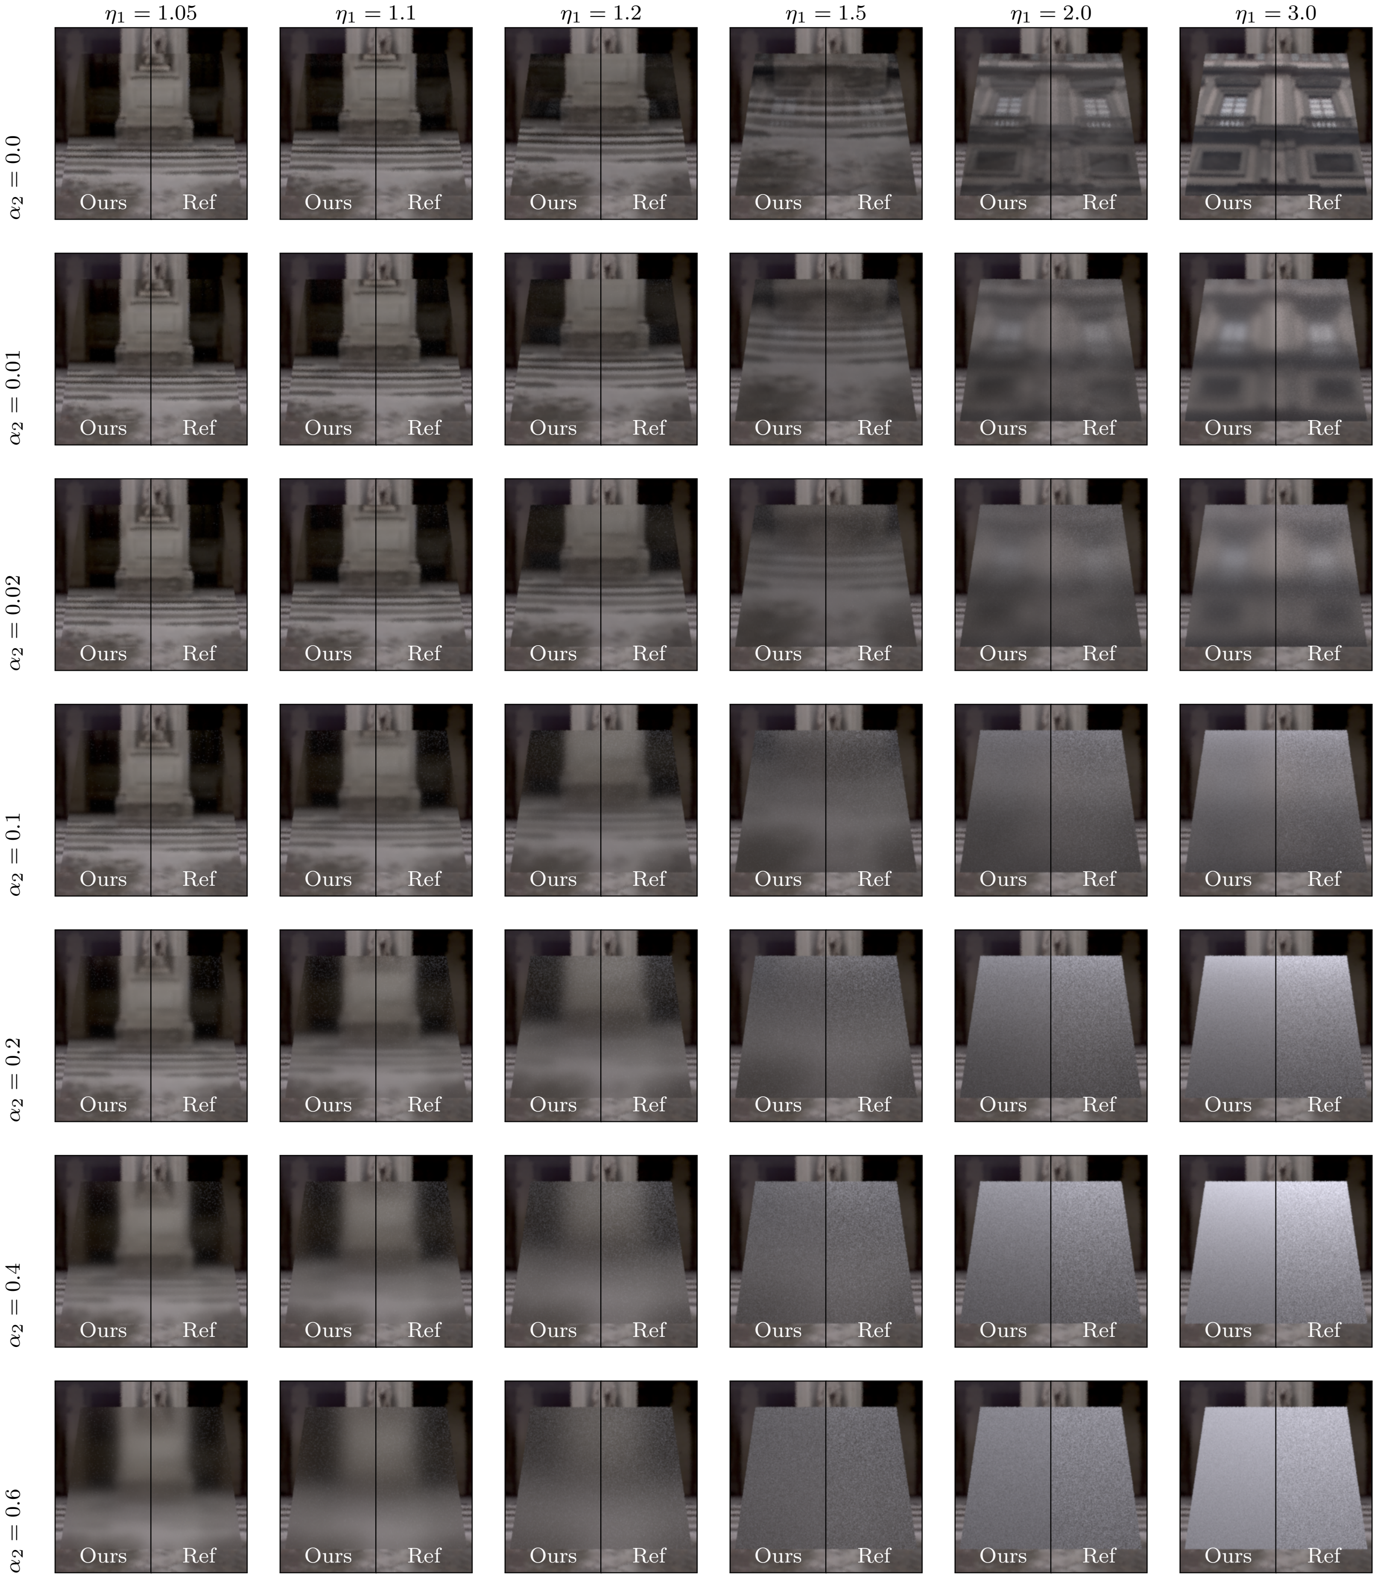

4 Scattering

We validated the ability of our method to handle participating media, and more specificaly scattering. In this scene, we used a sphere with a two layers materials: one layer consisting of a participating media with depth $d = 1$, absorption cross-section $\sigma_a = 0$, and varying scattering cross-section σ_s, and phase function anisotropy g. We display side-by-side our layered BSDF track two lobe at a time and a reference stochastic layered BSDF model accounting only for single scattering. Our model is a close match for the overall apparence. However, note that we cannot replicate the heavy-tails of the interaction between the HG phase function and the GGX microfacet model as highlighted by row $\sigma_s = 0.01$.

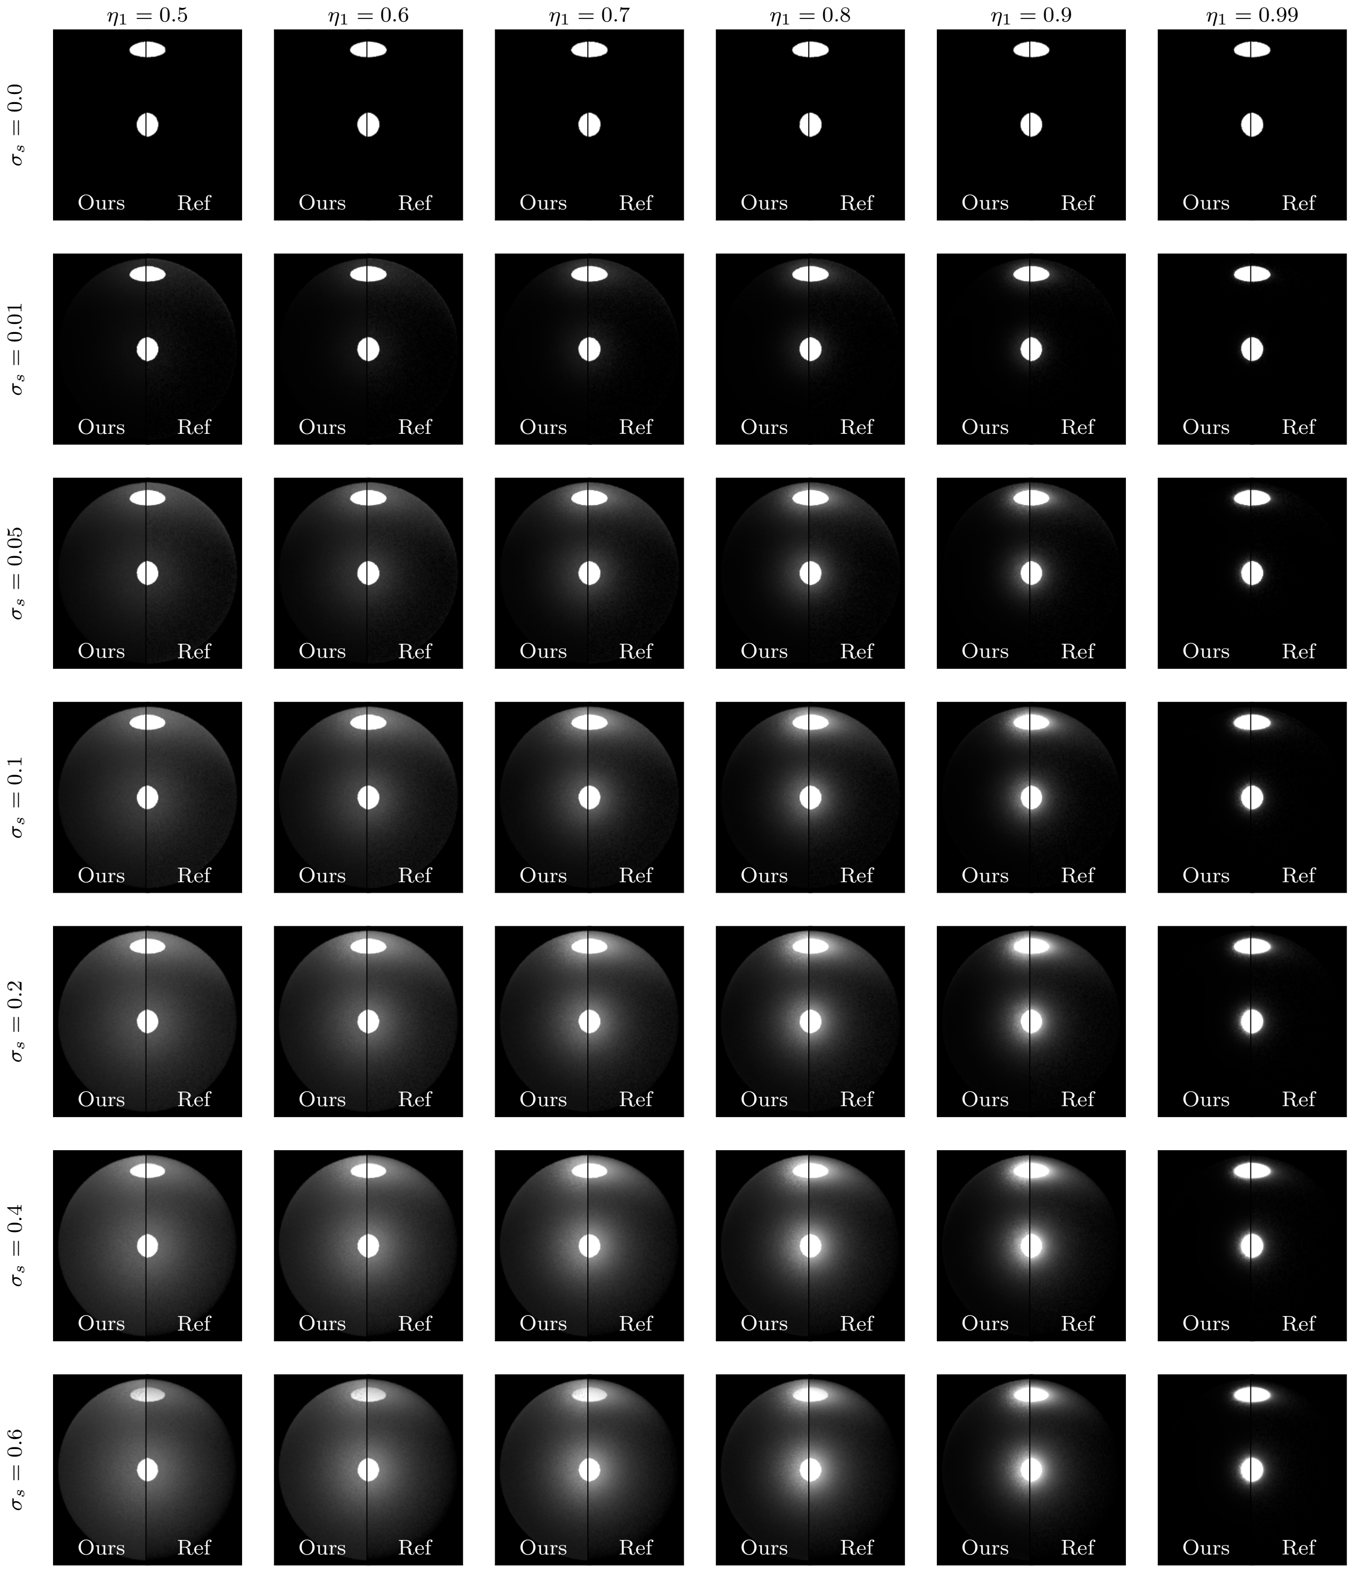

We also validated the ability of our method to match the reference when rendering participating media of increasing optical depth (using $h \in [0, 0.4]$) with an environement map lighting. In this scene, we use a pure mirror as the base layer on top of which a medium of increasing depth is applied. We specified $\sigma_a = 0$, $\sigma_s = 1$ and, $g = 0.9$ as constants for the medium. As shown below, our method faithfully reproduce the stochastic reference in such context.

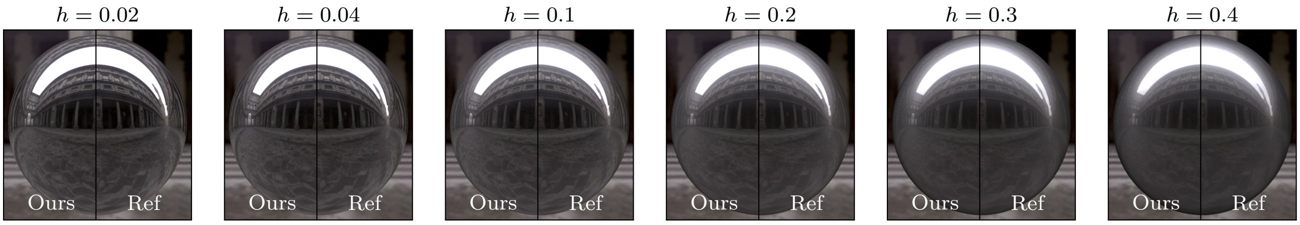

5 TIR Factor

We validated the impact of our TIR factor using a simple unit test. We set a coated sphere lit by two disk lights (one for the normal incidence, another for grazing angles) and run our algorithm with or without our TIR factor. We compare the different outputs with our Ground Truth code. We set the layered material to have two layers, a first smooth dielectric layer with varying index of refraction and a rough conducting base with $\boldsymbol{\eta} = 1 + [1, 0.5, 0.5]\,\mathbf{i}$, and $\alpha = 0.6$. As reported in the table, the TIR factor permits to avoid overestimating the contribution of short paths (such as TRT in this case). Not that even if we do not use the TIR factor, our layered BSDF is by construction energy conserving and will not leak energy.

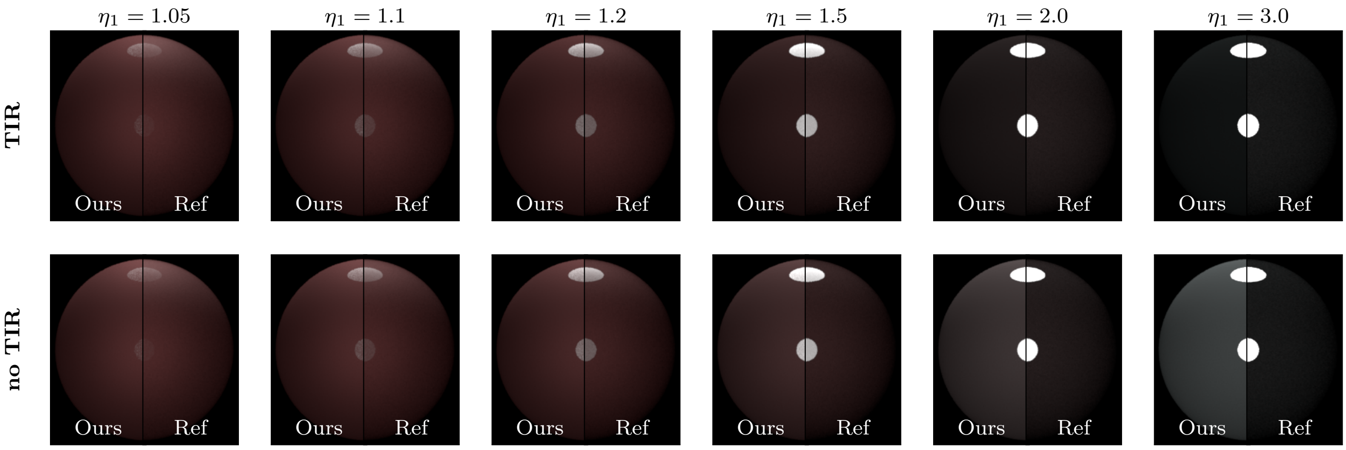

6 Comparison with Ground Truth

We compared our Forward and Symmetric models to our reference implementation. While both method provide similar apparences, it is to note that the Forward model overestimate the Fresnel term at grazing angles. This is in part due to the use of the average Fresnel. Just below, we add to the comparison the model of Weidlich and Wilkie [2007]. Because this model fails to correctly propagate the effect of the upper layer to the bottom layer and does not accounts for multiple scattering between layers, it cannot faithfully reproduce the different apparences. We also add the model of Jakob et al. [2014] using their LayerLab implementation and provide the storage cost and computation time required to generate the BSDF files. Since this implementation computes the transport on Beckmann microfacet models, we cannot match the reference renderings. But we can see all the properties we are interesting in (energy conservation, color saturation, etc). However, LayerLab has a non constant memory footprint and computation time. More precisely, its cost explodes for the type of appareances that are interesing for clear-coated shaders that are usually the main use in production of layered materials. In our tested models, it can take up to 2.4GiB and 22min to generate a BSDF file.

Metal Foil | Metallic Paint | Frosted Metal | Rough on Rough | Gold Coated

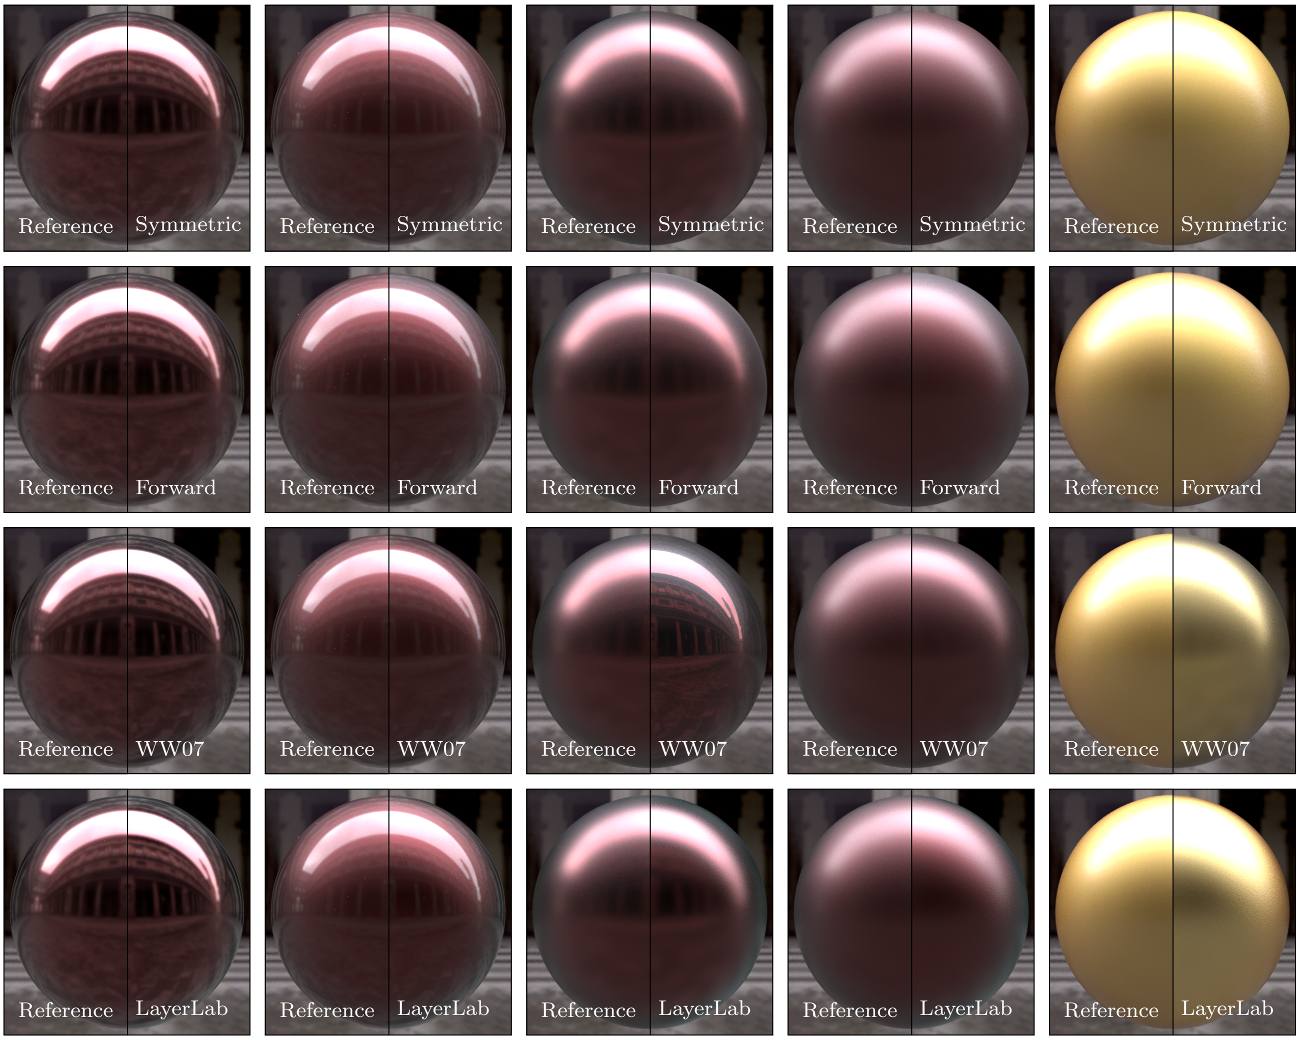

2.4GiB - 22min | 1.7GiB - 21min | 40MB - 10sec | 6.8MB - 2sec | 6.7MB - 4sec

7 The Roadster Scene

The ROADSTER scene illustrate how interesting layered materials are to depict the appearance of car paints. In this example, we used threee layered materials to cover the asset. We use the complex refractive index of cooper $\eta = [0.27 + 3.6\,\boldsymbol{i}, 0.98 + 2.37\,\boldsymbol{i}, 1.33 + 2.3\,\boldsymbol{i}]$ and coated it with a rough dielectric $\eta = 2,\ \alpha = 0.1$ layer. We also used a red rough conductor $\eta = [1 + \boldsymbol{i}, 1 + 0.1\,\boldsymbol{i}, 1 + 0.1\,\boldsymbol{i}]\,,\ \alpha = 0.1$ coated with a smooth dielectric $\eta = 1.5$ layer.