Supplementary Material of Many-Worlds Inverse Rendering

ZIYI ZHANG, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland
NICOLAS ROUSSEL, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland
WENZEL JAKOB, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

CCS Concepts: • **Computing methodologies** → **Rendering**.

Additional Key Words and Phrases: differentiable rendering

ACM Reference Format:
Ziyi Zhang, Nicolas Roussel, and Wenzel Jakob. 2025. Supplementary Material of Many-Worlds Inverse Rendering. *ACM Trans. Graph.* 1, 1 (September 2025), 4 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 Additional results on volume-based reconstruction

We present additional results on the exponential volume reconstruction experiment (see Figure 12 in the main document).

1.1 Microflake volume optimization

The results presented in the main document show the outcome of volume optimization after 200 iterations. Figure 1 further demonstrates that after 500 optimization iterations, the microflake volume successfully reproduces the surface appearance. The isotropic volume does not improve with more iterations.

Note that this is a purely volumetric reconstruction: The scene must be rendered with volumetric path tracer that interacts with phase functions instead of BRDFs, so a surface cannot be directly extracted from this result.

Fig. 1. **Converged microflake volumes.** We visualize the result after 500 iterations of optimization. An anisotropic microflake volume can mimic surface appearance. Increasing the *samples per pixel* (spp) during training slightly improves reconstruction quality at the expense of longer computation times.

1.2 SDF-parameterized microflake volume

In *radiance field reconstruction*, a common strategy to extract a surface from a volume is to parameterize the volume using a *signed distance function* (SDF) [Yariv et al. 2021; Wang et al. 2021; Miller et al. 2024]. To the best of our knowledge, no previous work has applied this method to physically based rendering without approximations. So we implemented this technique in Mitsuba 3 [Jakob et al. 2022] to compare with our method.

Convergence. Consistent with prior work [Yariv et al. 2021; Wang et al. 2021; Miller et al. 2024], we enforce the SDF property by incorporating an Eikonal loss in the optimization objective:

$$\mathcal{L} = \mathcal{L}_{\text{render}} + w\mathcal{L}_{\text{Eikonal}}, \tag{1}$$

where w is the weight of the Eikonal loss.

Figure 2a shows that a higher Eikonal loss weight, while necessary to enforce the SDF property, slows down optimization. In this scene, a minimum weight of 0.01 is required to prevent divergence of the Eikonal loss (see Figure 2b).

An SDF-parameterized volume recontruction is considered converged when the standard deviation σ of its implicit function distribution becomes sufficiently small, effectively removing volume particles outside the SDF surface. We observed this behavior in our experiments (Figure 2c). Furthermore, Figure 2d visualizes how SDF values are mapped to volume densities using the Gaussian implicit function distribution (Equation (15) in Miller et al.'s work [2024]), shown at the scale of our scene.

Finally, Figure 2e shows the optimization states of 500 iterations with an Eikonal loss weight of 0.01, proving the convergence of this method.

Surface extraction. While SDF-parameterized volume reconstruction achieves convergence, extracting a surface from the dense volume necessitates an additional optimization step. Moreover, the extracted surface often fails to align precisely with the original volume, leading to a reduction in visual quality.

We demonstrate this with a simple scene in Figure 3. The camera faces a thick microflake volume (with an extinction coefficient of 1000) enclosed in a cube. Our objective is to extract a surface BRDF, defined over the same cube, that reproduces the volume's appearance. For the surface BRDF, we use the GGX model [Walter et al. 2007], and for the volume's phase function, we use the SGGX microflake model [Heitz et al. 2015]. The SGGX matrix is homogeneous across the volume and is isotropic with parameters $S_{xx} = 0.01$, $S_{yy} = 0.01$, $S_{zz} = 1$, $S_{xy} = 0$, $S_{xz} = 0$, $S_{yz} = 0$, where all coefficients follow the definition of Equation (10) in Heitz et al.'s work.

At first glance, the conversion might seem straightforward: the dense volume ensures that rays interact at the same point as in the surface case, and both microfacet and microflake models define

Authors' Contact Information: Ziyi Zhang, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, ziyi.zhang@epfl.ch; Nicolas Roussel, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, nicolas.roussel@epfl.ch; Wenzel Jakob, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, wenzel.jakob@epfl.ch.

ACM 1557-7368/2025/9-ART
https://doi.org/10.1145/nnnnnnn.nnnnnnn

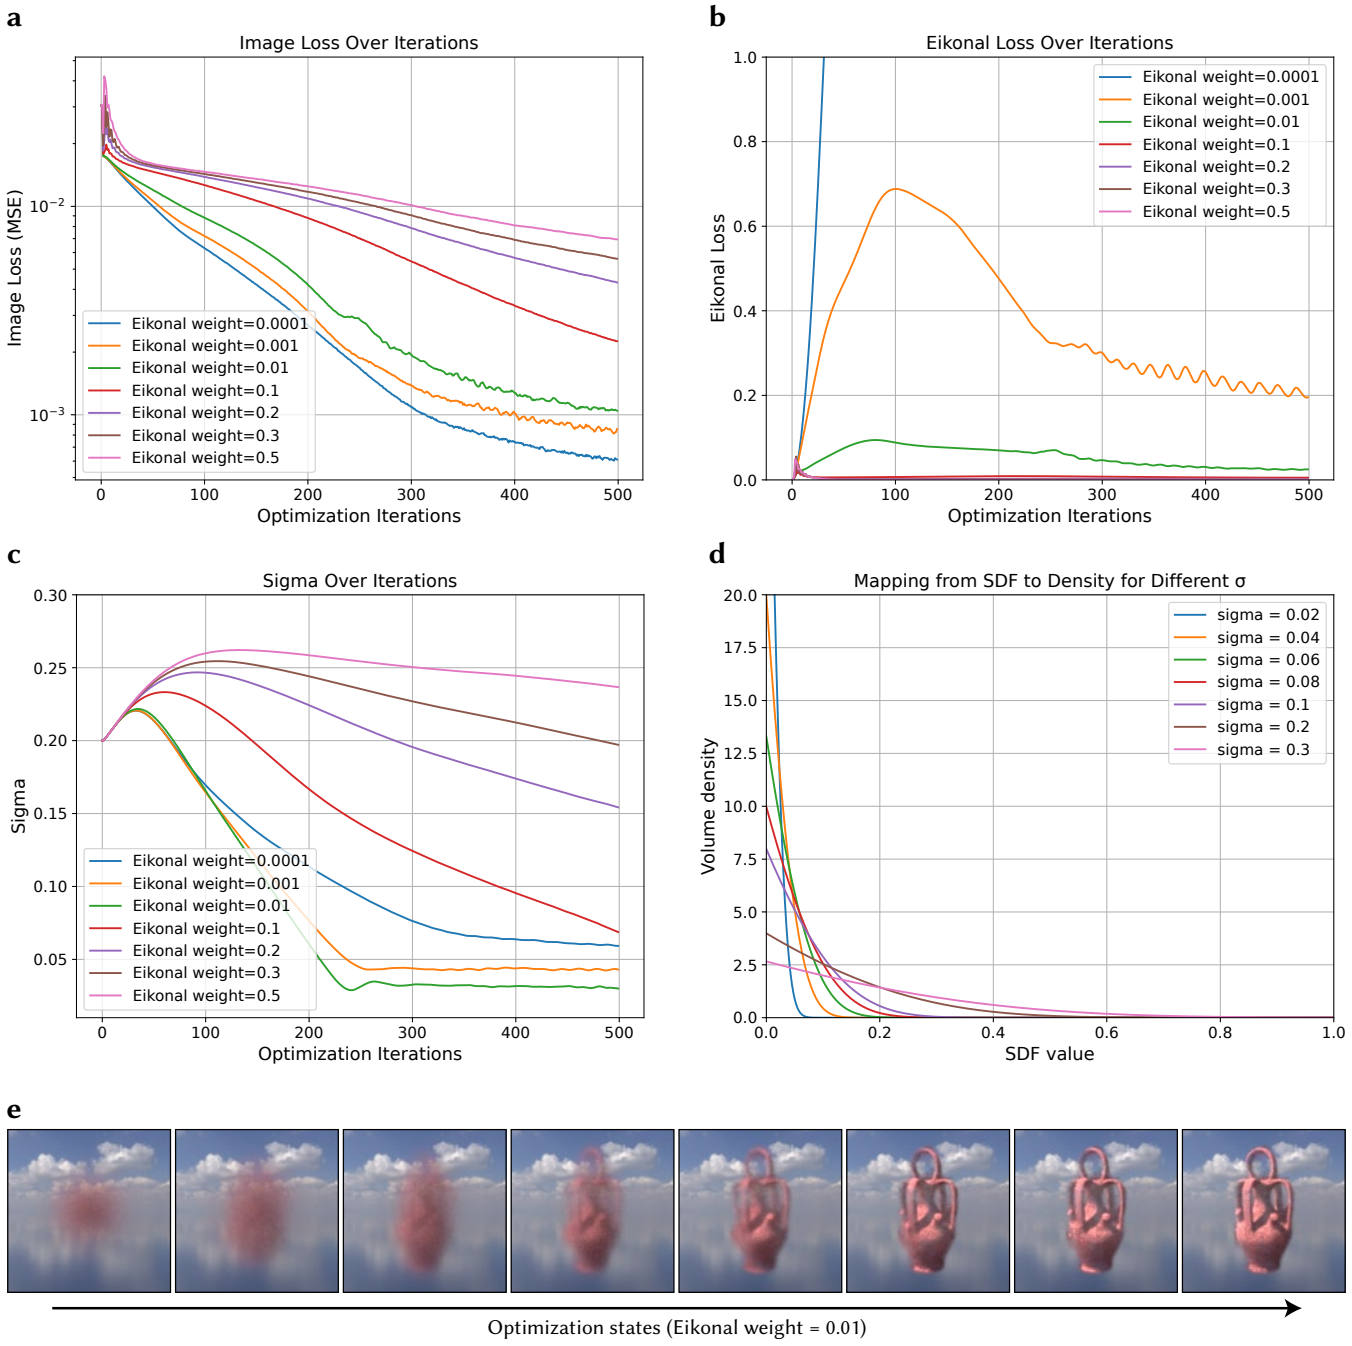

Fig. 2. **Optimization of SDF-parameterized microflake volumes.** We evaluate surface reconstruction using a volume parameterized by an SDF with an Eikonal constraint [Yariv et al. 2021; Wang et al. 2021]. **(a)** The image reconstruction loss over iterations reveals that stronger Eikonal regularization slows convergence. **(b)** A minimal Eikonal loss weight of 0.01 is needed to prevent divergence of the Eikonal loss in our scene. **(c)** The volume sharpens over time, as indicated by decreasing standard deviation σ of the implicit function [Miller et al. 2024]. **(d)** This sharpening is driven by the mapping from SDF values to densities and concentrates volume density near the surface to better approximate the reference surface. **(e)** Despite converging after 500 iterations with a stable Eikonal loss, the result remains inferior to our method in terms of surface quality and convergence speed. Overall, these results confirm that SDF-parameterized volume optimization can recover surface geometry, but it is significantly less efficient in both convergence speed and final quality.

their normal distributions (NDF) in a similar manner (the SGGX distribution is a symmetric GGX distribution). However, as shown in Figure 3b, the closest visual match occurs when the GGX roughness parameter (alpha) is set to 0.2120 — and even then, the match remains an approximation rather than an exact replication.

This mismatch is expected because the microfacet model is not an exponential model. Simply increasing the density of a microflake volume will not produce the same appearance of a microfacet surface. Previous work [Dupuy et al. 2016] has investigated the link between SGGX microflake models and GGX microfacet models, showing that the two can become equivalent when a *homogeneous* microflake volume is *semi-infinite* and the microflakes exhibit *non-symmetric* behavior when interacted with from the backside. This insight aligns with our observation that even with an additional stage of optimization, a perfect match between the two models is still unattainable.

In practice, the conversion is far more complex than this simple example. The optimized microflake properties are often *heterogeneous*. Additionally, the geometric normal (determined by the volume density) may not align with the NDF normal (determined by the microflake parameters), further complicating the conversion process. Consequently, optimizing a PBR volume to reconstruct a surface without approximations remains a difficult and open problem.

Speed. As shown in Figure 4, the computational cost per iteration is *one order of magnitude* higher than that of our approach, primarily because our method is surface-based and does not model shadowing or multiple scattering. The additional Eikonal loss also makes SDF-parameterized volume optimization slower than standard volume optimization.

Physically based rendering requires sampling entire light paths between sensors and emitters, a challenging task that makes volume-based methods inherently slower than surface-based ones. Two key factors contribute to this increased computational time:

- Multiple interactions: A light path interacts with the volume at several points. We need to sample these interaction positions in proportion to either the transmittance [Nimier-David et al. 2022] or the extinction-weighted transmittance.
- Iterative computation: As illustrated in Figure 5, connecting a shading point to a light source (for example, during next event estimation) requires only one shadow ray test in surface-based methods. In contrast, volume-based methods must also estimate the transmittance along the shadow ray using ratio tracking [Novák et al. 2014].

Together, these challenges make volume optimization significantly slower than surface optimization.

In *radiance surface reconstruction* (like NeuS [Wang et al. 2021]), we render a pixel by estimating a line integral along the camera ray through point sampling. The SDF parameterization enables efficient point sampling strategies — such as concentrating more samples near the SDF surface [Oechsle et al. 2021] — but it does not straightforwardly provide an easier solution to the aforementioned challenges in physical light transport.

Summary. Based on our experiments, we draw the following conclusions:

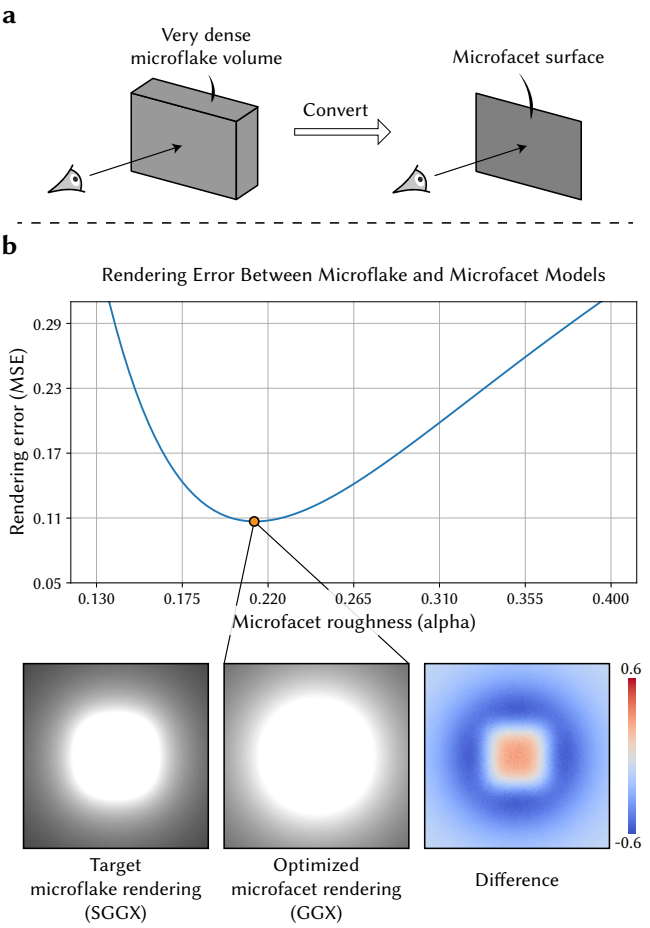

Fig. 3. **Extracting BRDFs from a simple volume. (a)** In a scene with an extremely dense microflake volume characterized by a *spatially homogeneous* phase function, our goal is to derive a surface BRDF that mimics the same appearance. **(b)** This task introduces a new optimization step since there is no analytical solution for the conversion. Furthermore, even in this trivial case, the conversion cannot exactly preserve the appearance of the original volume.

(1) **Surface Appearance**: SDF-parameterized microflake volumes can effectively mimic target surface appearances.
(2) **Surface Extraction**: Converting optimized volumes into surfaces with BRDFs remains an open challenge. Unlike radiance field reconstruction, where material properties and interreflection are absorbed into a radiance storage, the PBR problem requires a dedicated optimization step to extract a surface, and even then, the result is an approximation rather than an exact match.
(3) **Computational Efficiency**: Volumetric reconstruction is significantly slower than our method. In particular, the SDF parameterization does not alleviate the inherent computational costs of volume rendering.

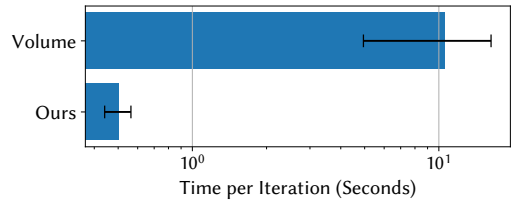

Fig. 4. **Average per-iteration computational cost.** Our approach achieves a 21.2× speedup in per-iteration computational cost compared to SDF-parameterized volume optimization. This comparison is based on identical hyperparameters (image resolution, samples per pixel) and both methods use the same texture storage in Mitsuba 3 [Jakob et al. 2022]. The reported times include all overheads, such as the primal rendering pass.

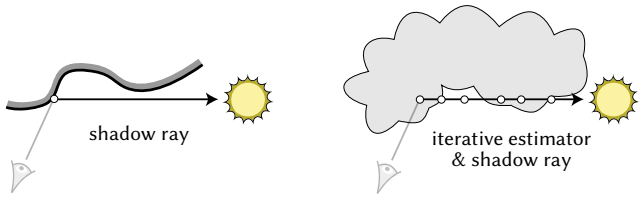

(**a**) Surface (**b**) Volume

Fig. 5. **Computational disparity between volume and surface rendering.** To connect a shading point to a light source (e.g., next event estimation), a surface renderer only requires one shadow ray test, while a volume renderer additionally need to estimate transmittance along the shadow ray. In a heterogeneous volume, this requires a costly iterative algorithm to estimate.

References

Jonathan Dupuy, Eric Heitz, and Eugene d'Eon. 2016. Additional Progress Towards the Unification of Microfacet and Microflake Theories. In *EGSR (EI&I)*. 55–63.

Eric Heitz, Jonathan Dupuy, Cyril Crassin, and Carsten Dachsbacher. 2015. The SGGX microflake distribution. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–11.

Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. 2022. *Mitsuba 3 renderer*. https://mitsuba-renderer.org.

Bailey Miller, Hanyu Chen, Alice Lai, and Ioannis Gkioulekas. 2024. Objects as Volumes: A Stochastic Geometry View of Opaque Solids. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 87–97.

Merlin Nimier-David, Thomas Müller, Alexander Keller, and Wenzel Jakob. 2022. Unbiased Inverse Volume Rendering with Differential Trackers. *ACM Trans. Graph.* 41, 4, Article 44 (July 2022), 20 pages. doi:10.1145/3528223.3530073

Jan Novák, Andrew Selle, and Wojciech Jarosz. 2014. Residual ratio tracking for estimating attenuation in participating media. *ACM Trans. Graph.* 33, 6, Article 179 (Nov. 2014), 11 pages. doi:10.1145/2661229.2661292

Michael Oechsle, Songyou Peng, and Andreas Geiger. 2021. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *Proceedings of the IEEE/CVF international conference on computer vision*. 5589–5599.

Bruce Walter, Stephen R Marschner, Hongsong Li, and Kenneth E Torrance. 2007. Microfacet Models for Refraction through Rough Surfaces. *Rendering techniques* 2007 (2007), 18th.

Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. 2021. NeuS: Learning Neural Implicit Surfaces by Volume Rendering for Multi-view Reconstruction. In *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (Eds.), Vol. 34. Curran Associates, Inc., 27171–27183. https://proceedings.neurips.cc/paper_files/paper/2021/file/e41e164f7485ec4a28741a2d0ea41c74-Paper.pdf

Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. 2021. Volume rendering of neural implicit surfaces. In *Thirty-Fifth Conference on Neural Information Processing Systems*.