Eurographics Symposium on Rendering 2025
B. Wang and A. Wilkie
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 44 (2025), Number 4

Supplemental Material: A Data-Driven Approach to Analytical Dwivedi Guiding

Darryl Gouder[1,2], Jiří Vorba[2], Marc Droske[2], Alexander Wilkie[1]

[1]Faculty of Mathematics and Physics, Charles University, Prague, Czech Republic
[2]Wētā FX, New Zealand

1. Sampled Slab Normal PDF proof

The analytical Dwivedi free-flight distance and scattering direction PDFs mentioned are conditioned on the slab normal used. In the case of Point Of Entry (PoE) and Closest Point (CP), the slab normal is deterministic. Therefore, the scattering direction PDF results in $p(\omega_z) = p_D(\omega_z|\mathbf{n_D})$. However *II*, and our techniques rely on a sampled slab normal, thus leading to a joint distribution:

$$p(\omega_z) = \int_\alpha p(\omega_z, \mathbf{n_D}) d\mathbf{n_D} \tag{1}$$

where α is the set of slab normals and n_D is the slab normal. This means the total density of some direction ω_o is the integral of the Dwivedi directional density over the set of slab normals. In other words, a fixed scattering direction's probability density depends on the slab normals. The joint distribution expands to

$$p(\omega_z, \mathbf{n_D}) = p(\omega_o|\mathbf{n_D}) \cdot p(\mathbf{n_D}) \tag{2}$$

Where $p(\omega_o|n_D)$ is the polar angle Dwivedi pdf and $p(\mathbf{n_D})$ is the density of selecting that particular slab normal. The probability density of sampling the slab normal, in our case, is another joint density - the probability mass of selecting the voxel and the probability mass of selecting the point, for BOPO, or the density of sampling the incident illumination direction, for BILL.

$$p(\mathbf{n_D}) = p(\mathbf{n_D}|\mathbf{v}) \cdot p(\mathbf{v}) \tag{3}$$

The full identity for $p(\omega_z)$ is:

$$p(\omega_z) = \sum_{i=0}^{j} p(v_j) \int_\mathbf{N} p(\omega_z|\mathbf{n_D}) \cdot p(\mathbf{n_D}|v_j) \cdot d\mathbf{n_D} \tag{4}$$

When a directional distribution is used for the sample space of slab normals, the integral is intractable. To avoid evaluating the integral, we propose approximating it using Stochastic MIS (SMIS) [WGGH20] to compute an estimate. As shown in their paper, the biased approximation of the density values will not introduce bias to the volume random walk estimator. A 1-sample estimator of $p(\omega_o)$ is:

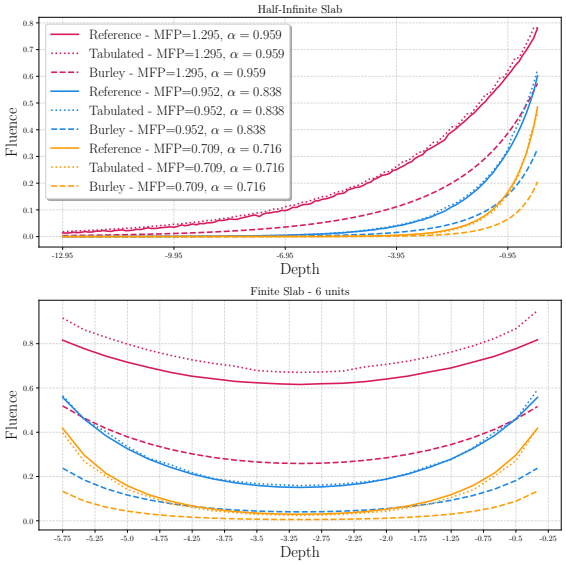

Figure 1: We brute-force-computed fluence measurements against area integration and a diffusion profile to compute the volume transport term at decreasing depths in a half-infinite slab (top row). We show that our profile maintains its accuracy as the fluence meter goes deeper into the volume, whilst the Burley Diffusion Profile significantly underestimates the fluence.

$$\langle p(\omega_z) \rangle = \frac{p(v_j) \cdot p(\omega_z|\mathbf{n_D}) \cdot p(\mathbf{n_D}|v_j)}{p(v_j) \cdot p(\mathbf{n_D}|v_j)} \tag{5}$$

and due to the cancellation, we have:

$$\langle p(\omega_z) \rangle = p_D(\omega_z \mid \mathbf{n_D}) \tag{6}$$

Note that what we have proven here is similarly applicable to BOPO and also to Incident Illumination.

The same issues also hold for the free-flight distance sampling PDF but are also handled equivalently.

2. Evaluating the Burley Diffusion Profile

We evaluate the Burley Diffusion profile for computing CDF weights by comparing it to a depth-aware tabulated diffusion profile. The Burley diffusion profile is fitted using surface integration only, neglecting transport between surface and volume points or points on opposite sides of a volume boundary. It can inadequately represent the energy transport between points that deviate from the original geometric configuration. The tabulated diffusion profile preserves the geometric assumptions of the semi-infinite medium while introducing an additional depth parameter. Our profile is simulated through extensive Monte Carlo random walks, originating at a fixed depth interval. The tables are parameterized by depth and radial distance. This parameterization enables a more accurate approximation of the transport operator between surface and volumetric points. The evaluation is composed of two parts. The first is a fluence measurement at various depths in the context of semi-infinite and finite slabs. The measured values at various depths showed that the Burley diffusion profile significantly underestimated the fluence compared to the tabulated and reference profile, which matched. The shapes of the curves, however, were very similar which significantly impacts CDF construction. When used as weight generators for normalized CDFs, two functions with comparable shapes yield similar CDFs despite magnitude differences, as normalization preserves only relative distribution patterns. The second experiment showed that when the CDFs were constructed using either the Burley Profile or the Tabulated profile, the CDFs were very similar.

However, we show that their shapes are similar While the semi-infinite planar assumption remains a limitation for arbitrary geometric setups, our approach substantially reduces error over the Burley profile when approximating volume point fluence in such scenarios.

2.1. Fluence Simulations

Figure 1 shows fluence measurements on a half-infinite slab and a finite slab of 6 scene units . In both cases, fluence differs only with increasing depth. The reference data was generated using Monte Carlo path tracing initiated below the surface, with continual measurements down to -12.95 units for the half-infinite slab. The volume parameters were $\sigma_t = (0.772, 1.05, 1.41)$ and $\alpha = (0.958, 0.838, 0.716)$.

We observe that the fluence diminishes with increasing depth from the surface. The tabulated and Burley profile measurements were obtained by sampling a surface point and estimating volume transport using the respective profiles, eliminating the need for stochastic volume random walks. As expected, the Burley profile consistently underestimates the fluence, whereas our proposed method demonstrates closer correspondence with the reference data. The finite slab configuration, with identical volumetric properties and uniform illumination from both surfaces, similarly demonstrates the Burley diffusion profile's tendency to underestimate fluence values.

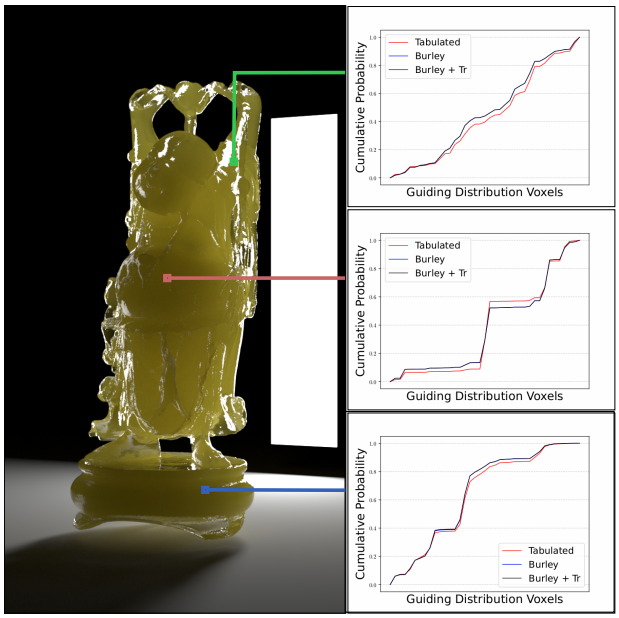

Figure 2: We show 3 pixels that are hand-picked and compare the CDFs using different diffusion profiles for weight construction.

2.2. Guiding Distribution CDF comparison using different profiles

In this section, we examine the CDFs generated using our tabulated diffusion profile versus the Burley profile. Figure 1 reveals that while the Burley profile consistently underestimates energy values, the shapes of both profiles exhibit some similarity. This shape similarity significantly impacts CDF construction. When used as weight generators for normalized CDFs, two functions with comparable shapes yield similar CDFs despite magnitude differences, as normalization preserves only relative distribution patterns. Figure 2 compares constructed CDFs at the volume entry point for hand-picked pixels using different weighting schemes and Figure 3 demonstrates this in equal-sample renders. In the rest of our results, we use the Burley diffusion profile.

3. Guiding Cache CDF Sizes

The size of the CDFs is directly related to the resolution of the guiding cache. Table 2 shows the number of fitted voxels, which in turn reflects the size of the CDFs. Given the scene and cache resolution, we list the CDF sizes. Without the CDF caching, constructing these CDFs for each random walk would be prohibitively expensive.

We show that unless the cache is very coarse, such as in the ENVMAP BUDDHA with $5 \times 5 \times 5$ cache, the boundary is adequately covered with distributions, be it boundary points or directional distributions. Figure 4, 5, 6 and 7 exhibit this.

4. Combining ours with Classical using MIS

We include the results for some of the other scenes not included in the main text in Figures 8 and 9.

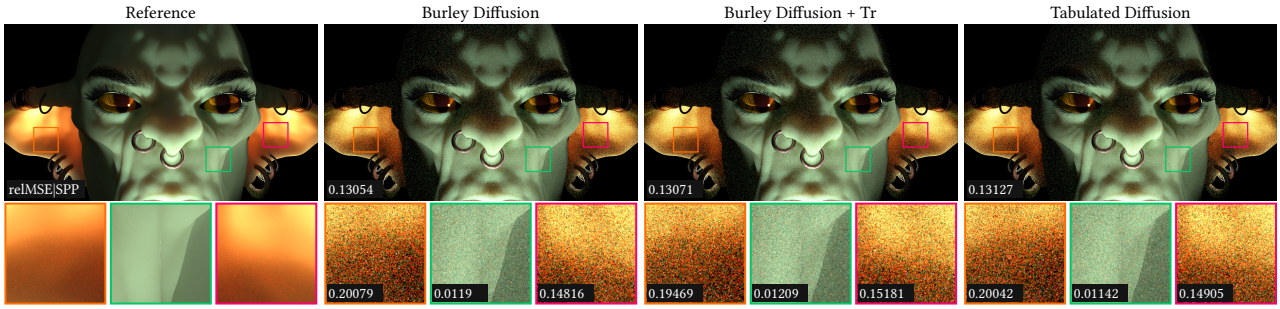

Figure 3: Comparing the different weighting strategies for CDF construction, where we can see that the variance is not affected by the weighting heuristic. Images were rendered using 32 samples per pixel, and render times were very similar.

5. Training and fitting timings

Our approach learns the radiance field at a volume boundary through a two-stage process: a pre-training step using Classical sampling followed by von Mises-Fisher (vMF) distribution fitting for the BILL algorithm. In contrast, BONO eliminates the fitting stage entirely while maintaining equivalent training render times. VPG presents the most complex training pipeline, requiring both an adjoint-driven rendering process for its guiding structure and an additional rendering pass to generate reference images for its Guided Russian Roulette. Given these fundamental differences in training architectures, direct timing comparisons would be misleading and provide little insight into algorithmic efficiency.

References

[WGGH20] WEST, REX, GEORGIEV, ILIYAN, GRUSON, ADRIEN, and HACHISUKA, TOSHIYA. "Continuous Multiple Importance Sampling". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 39.4 (July 2020). DOI: 10.1145/3386569.3392436 1.

Figure 4: ENVMAP BUDDHA Guiding cache size resolution comparison

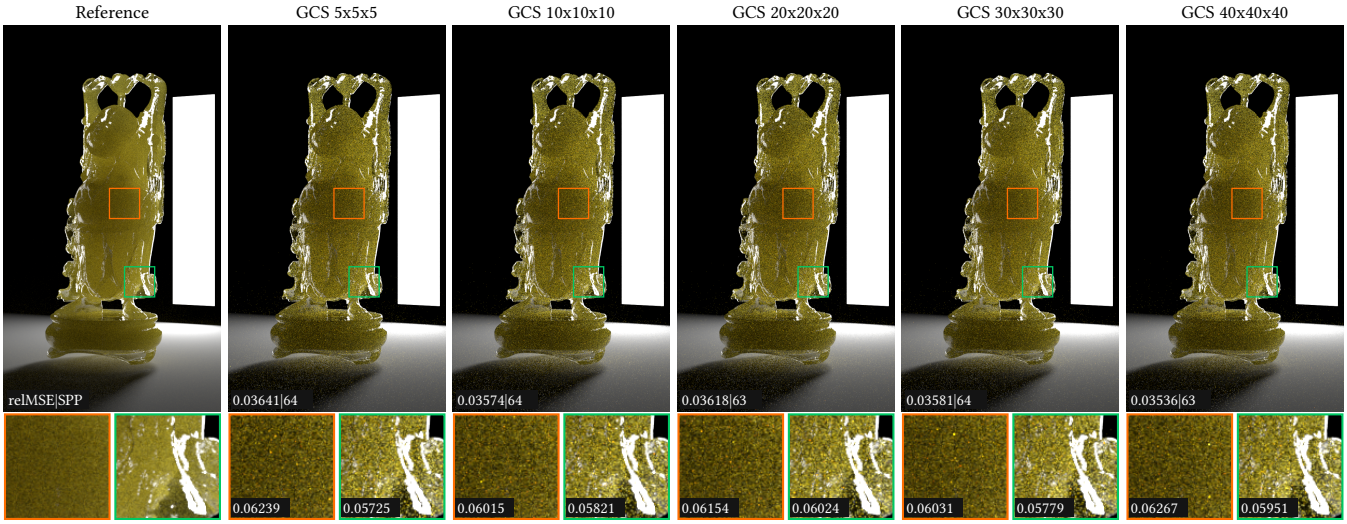

Figure 5: BACKLIT BUDDHA Guiding cache size resolution comparison

	Guiding Cache Size				
Scene	5x5x5	10x10x10	20x20x20	30x30x30	40x40x40
ENVMAP BUDDHA	94	445	2108	4949	9003
BACKLIT BUDDHA	5	38	139	327	488
DAVID	25	101	268	440	613
ORC	19	81	322	741	1253

Table 1: For each guiding cache configuration, we show the average size of the CDF (which is the number of fitted voxels).

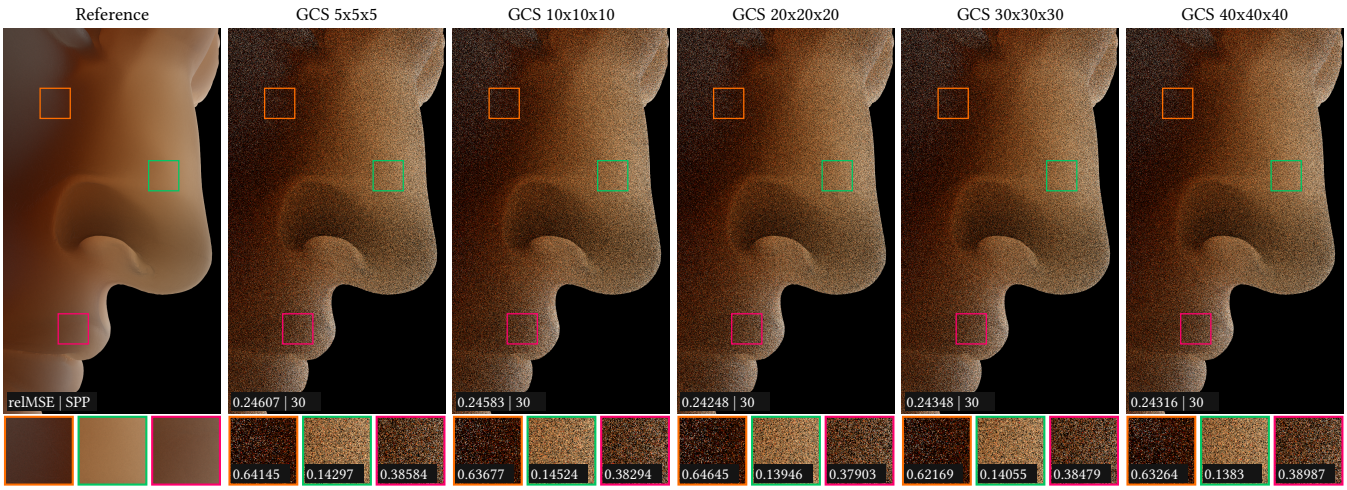

Figure 6: DAVID BUST Guiding cache size resolution comparison

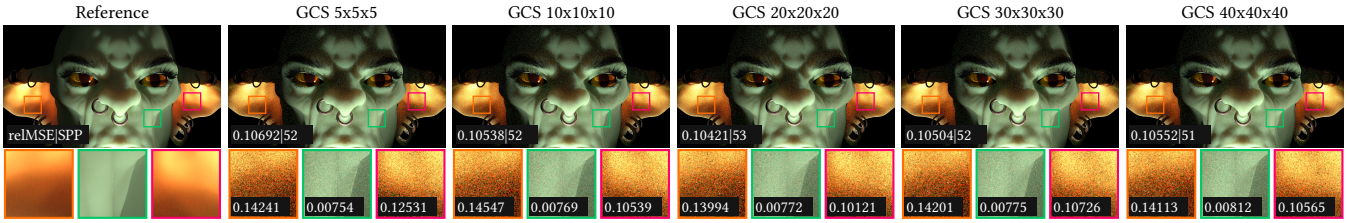

Figure 7: ORC Guiding cache size resolution comparison

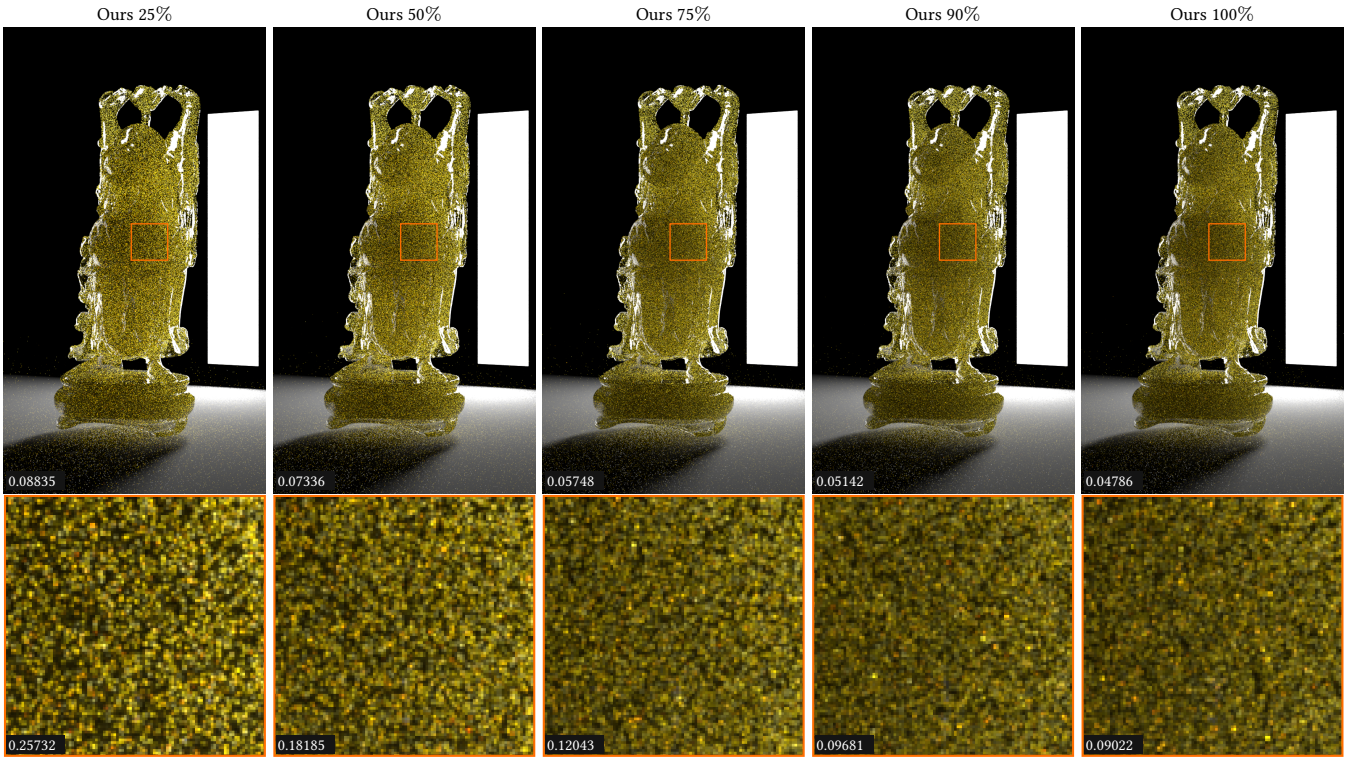

Figure 8: Buddha MIS comparisons

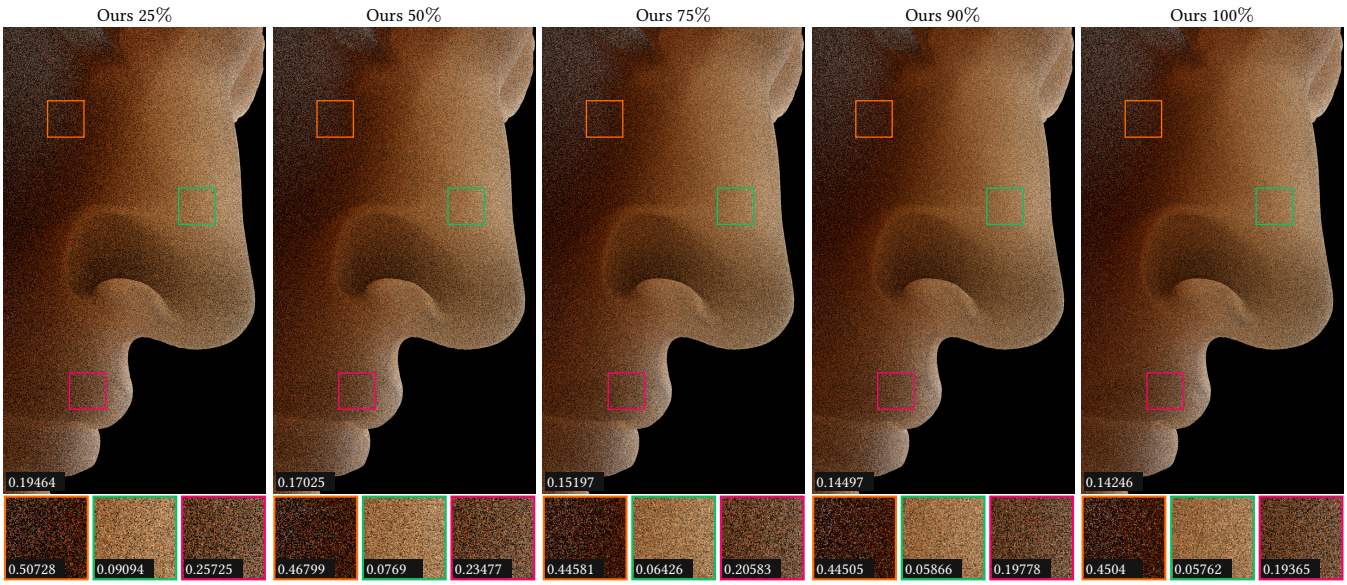

Figure 9: David Bust MIS comparisons

Scene	BTDF	Mean Free Path (scene units)	Single Scattering Albedo
Envmap Buddha	Rough Dielectric($\alpha = 0.3, \eta = 1.3$)	[0.005, 0.005, 0.005]	[0.8, 0.2, 0.75]
Backlit Buddha	Smooth Dielectric (n = 1.2)	[0.13, 0.09, 0.0671]	[0.959, 0.764, 0.678]
David	Diffuse Transmitter ($\rho = [0.8, 0.8, 0.8]$)	[0.0133, 0.011, 0.009]	[0.983, 0.926, 0.878]
Orc	Diffuse Transmitter ($\rho = [0.7, 0.9, 0.7]$)	[0.009, 0.006, 0.005]	[0.988, 0.958, 0.927]

Table 2: Material and volume parameters for the scenes we evaluated. Triplets are

Scene	Envmap Buddha	Backlit Buddha	David	Orc
Rendering	52.95	8.654	35.404	181.753
vMF Fitting	12.7	0.016	0.046	0.747
Total	65.65	8.67	35.45	182.5

Table 3: Training times for all the scenes, in seconds.