Eurographics Symposium on Rendering 2026
I. Gkioulekas and A. Jarabo
(Guest Editors)

Volume 45 (2026), Number 4

A Hybrid Neural-Microfacet BRDF Model for Real-Time Rendering: Supplemental Material

L. De Oliveira[1,2], A. Karpova[1], G. Nader[1], A. Houdard[1], P. Mézières[2], D. Rioux-Lavoie[1], R. Pacanowski[2]

[1]Ubisoft La Forge, [2]Inria

Complementary to the companion video and the GLSL implementation of our hybrid model, this supplemental document provides additional theoretical details of our implementation (Section 1). It also provides additional results that illustrate the editing capabilities of our model (Section 3), and additional comparisons of the error in BRDF space (Section 5) and image space (Section 4). Section 7 provides additional results for the importance sampling. Finally, Section 8 gives more details on the path-traced scene shown in the main paper.

1. Details on our implementation

As introduced in the main paper, our hybrid model consists of two components. The first is the analytical one f_a implemented as a microfacet-based BRDF model, and the second one is the neural one ψ.

Analytical BRDF Model Component As described in the main paper, the analytical BRDF we used is based on the microfacet model as expressed by [WMLT07], with a GGX distribution term:

$$f_a(\omega_i, \omega_o; p) = \frac{k_d}{\pi} + \frac{F(\omega_i, \omega_o, \eta) G(\omega_i, \omega_o, \alpha) D(\omega_i, \omega_o, \alpha)}{4(\omega_i \cdot n)(\omega_o \cdot n)}, \quad (1)$$

where n is the surface normal. The first models the diffuse contribution where k_d is the diffuse albedo. The second term models the specular reflection with three key components: the normal distribution D, the geometry term G, and the Fresnel term F. The normal distribution describes the statistical distribution of the microfacet orientations. The GGX distribution [WMLT07] is defined as:

$$D(\omega_i, \omega_o, \alpha) = \frac{\alpha^2}{\pi((n \cdot h)^2(\alpha^2 - 1) + 1)^2} \quad (2)$$

where $\mathbf{h} = \frac{\omega_i + \omega_o}{\|\omega_i + \omega_o\|}$ is the half-vector. The normal distribution is can be made anisotropic [Hei14] by setting two roughnesses α_x and α_y (aligned on the axes $\mathbf{x}$ and $\mathbf{y}$):

$$D_{aniso}(\omega_i, \omega_o, \alpha) = \frac{1}{\pi \alpha_x \alpha_y} \frac{1}{\left(\frac{(\mathbf{x} \cdot \mathbf{h})^2}{\alpha_x^2} + \frac{(\mathbf{y} \cdot \mathbf{h})^2}{\alpha_y^2} + (n \cdot \mathbf{h})^2 \right)^2} \quad (3)$$

The geometry term G accounts for shadowing and masking effects. It is computed as the product of two masking functions $G(\omega_i, \omega_o, \alpha) = G_1(\omega_i, \alpha) G_1(\omega_o, \alpha)$ where $G1$ is defined as:

$$G_1(\omega, \alpha) = \frac{2(n \cdot \omega)}{(n \cdot \omega) + \sqrt{\alpha^2 + (1 - \alpha^2)(n \cdot \omega)^2}} \quad (4)$$

Finally, the Fresnel term F describes how reflectance varies with angle between the outgoing direction and the half-vector, depending on the material's index of refraction η:

$$F(\omega_i, \omega_o, \eta) = \frac{(g - (\omega_o \cdot \mathbf{h}))^2}{2(g + (\omega_o \cdot \mathbf{h}))^2} \left(1 + \frac{((\omega_o \cdot \mathbf{h})(g + (\omega_o \cdot \mathbf{h})) - 1)^2}{((\omega_o \cdot \mathbf{h})(g - (\omega_o \cdot \mathbf{h})) - 1)^2} \right) \quad (5)$$

where $g = \sqrt{\eta^2 + (\omega_o \cdot \mathbf{h})^2 - 1}$ and assuming that the IOR of the ambient medium is equal to one.

Note that other analytical models could have been used. However, this formulation strikes a good compromise between physical plausibility and ease of optimization. More complex models, such as Disney's BRDF [Bur12], could represent a wider range of appearances, but since they have more parameters, we observed numerical instabilities during optimization.

Neural Component Our MLPs use the "HardGELU" activation as introduced in [VSW*23] rather than the common "ReLU", this is an empirically better choice for shallow MLPs, resulting in a more complex nonlinearity for a cheap evaluation cost.

$$\text{HardGELU}(x) = \begin{cases} 0, & \text{if } x < -\frac{3}{2}, \\ x, & \text{if } x > \frac{3}{2}, \\ \frac{x}{3}\left(x + \frac{3}{2}\right), & \text{otherwise.} \end{cases}$$

2. Datasets and training

We tested our approach on three datasets:

- MERL [MPBM03], which consists of 100 measured materials, which we separate into 5 families (fabric, dielectric, metal, and phenolic) to ease the reading of the different results. The provided measurements (90×90×180 directions) are expressed using the Rusinkiewicz [Rus98] 3D parametrization modified with a non-linear mapping of the zenithal angle, θ_h, of the half-vector $\mathbf{h} = (\theta_h, \phi_h)$.

- UTIA [FV14], which consists of 150 anisotropic materials. The angular resolution of the measurements is 115 degrees for the elevation angle and 7.5 degrees for the azimuthal angle. The angles are expressed in the spherical parametrization.

- RGL [DJ18], consists of 62 materials, which we separate into three families: fabric, dielectric, and metal. The angular resolution is adaptive with respect to the material frequency.

Remember that the RGL database is model-driven since the provided directions have been sampled according to the chosen underlying microfacet model. In contrast, the MERL and UTIA databases are model-agnostic and are measurement-oriented.

To evaluate our fitting results in terms of BRDF space, we compute the Symmetric Mean Absolute Percentage Error (SMAPE) by sampling both the target and the estimated BRDF with 10^6 directions $\mathcal{B} = \{(\omega_\mathbf{i}, \omega_\mathbf{o})\}$.

$$SMAPE(f,\hat{f}) = \frac{2}{10^6} \sum_{\omega \in \mathcal{B}} \frac{|f(\omega) - \hat{f}(\omega)|}{|f(\omega)| + |\hat{f}(\omega)|}. \tag{6}$$

Where f is the reference measured BRDF and $\hat{f}$ the one obtained with the model.

To evaluate our rendering results in image space, we compute the Mean Average Error (MAE) between the rendering of the reference measured BRDF and the modeled one. These renderings are computed with 1024 samples per pixel (spp) in a simple test scene consisting of a sphere lit by the `Uffizi` environment map. The MAE is then computed for all pixels and color channels as follows:

$$MAE(f,\hat{f}) = \frac{1}{N*3} \sum_{c=1}^{3} \sum_{i=1}^{N} |f_c[i] - \hat{f}_c[i]|. \tag{7}$$

Where $f_c[i]$ (resp. $\hat{f}_c[i]$ is the i-th pixel of the c-th color channel of the rendered image of the BRDF f (resp. $\hat{f}$).

To be complete, we also provide results using the widely-used Peak Signal to Noise Ratio (PSNR) image-to-image metric. This quantity is evaluated as follows:

$$PSNR(f,\hat{f}) = 10\log_{10}\left(\frac{\max(f)^2}{MSE(f,\hat{f})}\right), \tag{8}$$

where

$$MSE(f,\hat{f}) = \frac{1}{N*3} \sum_{c=1}^{3} \sum_{i=1}^{N} (f_c[i] - \hat{f}_c[i])^2. \tag{9}$$

3. Additional Editability and Interpolation Examples

We give additional examples of editability on the different datasets: MERL (Figure 1), RGL (Figure 2) and UTIA (Figure 3). Figure 4 shows an example of editing an initially isotropic material to make it anisotropic by adding and modifying the roughness parameters. We also give additional examples of material interpolations in Figure 5.

4. Additional Error Metrics in Image Space

In this section, we show additional renderings to compare our model and the Neural model from [ZRW*24] at comparable network size and similar number of parameters. We also add, for comparison, the analytical component of our model. Table 1 gives the reconstruction error values in terms of PSNR (8). It shows that at comparable architecture, our model consistently outperforms the neural model. Figures 6 to 12 present these comparisons on rendered materials from the three tested datasets (see Section 2). For each rendering, we also add the MAE metric for the results obtained. Figures 6 and 7 are on materials from the MERL dataset, figures 8 and 9 are from the isotropic RGL dataset, figures 10 and 11 are from the anisotropic RGL dataset and figure 12 is from the UTIA dataset.

All images from Figure 6 to 12 are computed at a resolution of 512×512 pixels using our one-bounce path-tracer implemented in PyTorch.

Table 1: PSNR comparisons in Image space, between our model and the neural one, across the 3 tested datasets and for two sets of parameters (12 and 16). For each dataset and each 3-layer MLP configuration, we show the average PSNR (higher values with a green background are the best) for all materials of the corresponding dataset. Our Hybrid model outperforms the neural one at equal capacity in most cases except the 64x3 MLP with $l = 4$ for the RGL dataset.

MLP	12 params		16 params	
	Ours l=4	Neural l=12	Ours l=8	Neural l=16
MERL (GGX PSNR: 32.8935)				
16x3	39.3681	35.7522	39.7280	35.7744
32x3	40.5024	38.7944	40.6165	37.9961
64x3	41.1915	40.7123	41.5204	40.9319
RGL (GGX PSNR: 33.5113)				
16x3	35.5394	34.7826	36.1161	34.4680
32x3	37.5375	36.0590	37.5778	36.4861
64x3	39.1842	39.4303	39.0731	38.9614
UTIA (GGX PSNR: 32.5828)				
16x3	38.1736	36.4343	38.4437	36.8398
32x3	38.9011	38.3525	39.1572	37.4699
64x3	39.6182	39.1176	39.7517	38.9291

5. In-depth Error Analysis in BRDF Space

Figures 13 to 21 show the SMAPE value for each BRDF for GGX models, our model, and the Neural model from [ZRW*24] for same number of parameters. For sake of readability we regroup materials from the same dataset in the same figure and we sorted the result in decreasing SMAPE order for our method. We also separated the figures for each architecture size: Figures 13 to 15 present the fitting results using the SMAPE metric for an MLP of 16x3. Figures 16 to 18 present the fitting results using the SMAPE metric for an MLP of 32x3. Figures 19 to 21 present the fitting results using the SMAPE metric for an MLP of 64x3. Finally, we report the averaged error numbers for the SMAPE metric in Tables 2 to 4, that shows that our model consistently outperform the neural model also in term of BRDF reconstruction metric.

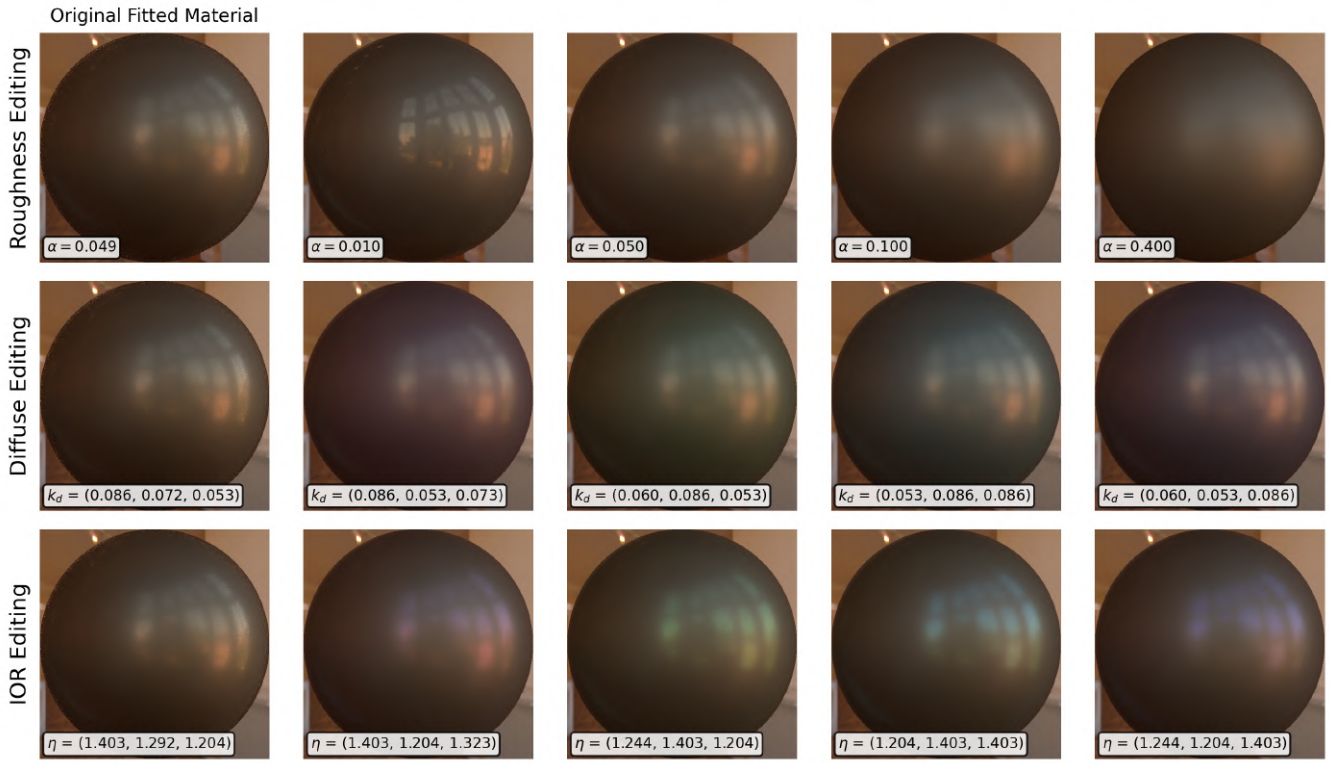

Figure 1: Editing a MERL BRDF (`alum-bronze`)

Table 2: SMAPE comparisons in BRDF space on the MERL dataset, between our model and the neural one, we show the average SMAPE (lower values with a green background are the best) for all materials in that dataset.

MERL (GGX SMAPE: 0.36)

MLP	12 params		16 params	
	Ours l=4	Neural l=12	Ours l=8	Neural l=16
16x2	0.2815	0.6006	0.2569	0.4869
16x3	0.2389	0.3890	0.2331	0.3921
32x2	0.2470	0.3524	0.2126	0.4055
32x3	0.2003	0.2537	0.1852	0.2866
64x2	0.2174	0.2681	0.1782	0.2400
64x3	0.1800	0.2085	0.1686	0.2072

6. Sparse training renders

Figures 22 and 23 shows the rendering results for BRDFs that were held out during training on half the MERL dataset. For each BRDF z and p were optimized separately with the MLP weights frozen.

7. Additional importance sampling results

Figure 24 and 25 show that using only GGX importance sampling is sufficient to achieve good convergence and yields results close to the reference obtained from a tabulated BRDF (similar

Table 3: Same metrics as computed in Table 2, on the RGL dataset.

RGL (GGX SMAPE: 0.64)

MLP	12 params		16 params	
	Ours l=4	Neural l=12	Ours l=8	Neural l=16
16x2	0.6184	0.7552	0.5977	0.7202
16x3	0.6095	0.6669	0.5777	0.6697
32x2	0.5899	0.6293	0.5855	0.6449
32x3	0.5687	0.6549	0.5491	0.6041
64x2	0.5925	0.5942	0.5785	0.6082
64x3	0.5100	0.5376	0.5017	0.5564

Table 4: Same metrics as computed in Table 2, on the UTIA dataset.

UTIA (GGX SMAPE: 0.19)

MLP	12 params		16 params	
	Ours l=4	Neural l=12	Ours l=8	Neural l=16
16x2	0.0958	0.1173	0.0963	0.1088
16x3	0.0913	0.1139	0.0855	0.0998
32x2	0.0778	0.1046	0.0733	0.1009
32x3	0.0677	0.0833	0.0631	0.0868
64x2	0.0634	0.0878	0.0561	0.0609
64x3	0.0483	0.0657	0.0442	0.0686

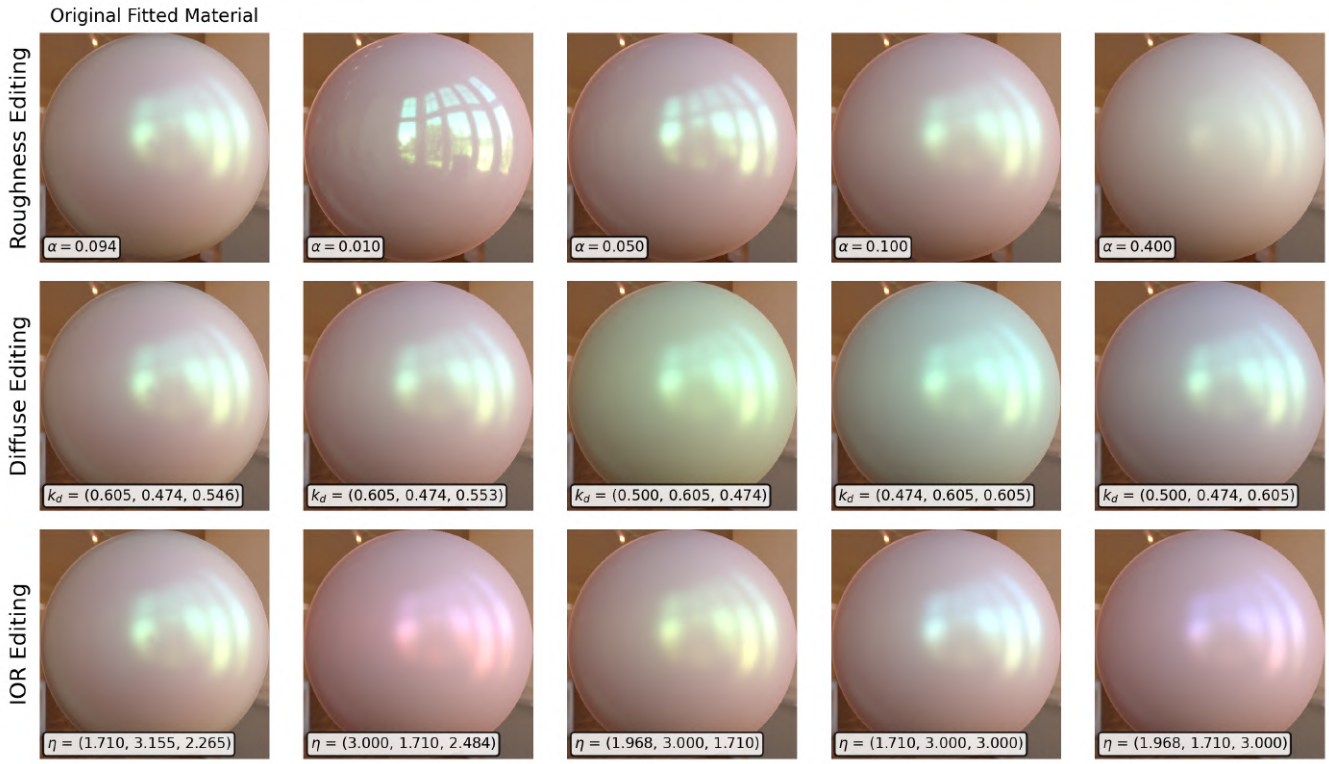

Figure 2: Editing a RGL BRDF (`aurora_white_rgb`)

to [LRR04a]). Figure 26 shows renderings obtained with GGX importance sampling at different samples per pixel levels.

8. Details on path-traced scene

Figure 27 illustrates some rendered images computed using global illumination implemented on our GPU path-tracer.

3D models The initial scene (the living room) was taken from blendswap, to which we added the xyz-dragon (Stanford 3D repository), a teapot [LRR04a], and the fertility statue (AIM@SHAPE Repository).

Materials We selected different materials from the different databases as follows:

- From UTIA [FV14]:
  - Wall: `m134_wallpaper03`
  - Wall Edge: `m124_plaster01`
  - Floor: `m150_wood63`
  - Sofa: `m009_corduroy01`
  - Pillow: `m071_fabric106`, with an edited k_d
- From MERL [MPBM03]:
  - Teapot handle : `neoprene-rubber`
  - Dragon : `yellow-plastic`
  - Fertility : `white-marble`
- From RGL [DJ18]:
  - Coffee Table : `cm_white`
  - Teapot body : `aniso_copper_sheet`
  - Left Vase : `aniso_miro_7`
  - Middle Vase : `cc_amber_citrine`
  - Right Vase : `cc_nothern_aurora`

References

[Bur12] BURLEY B.: Physically-Based Shading at Disney. 1

[DJ18] DUPUY J., JAKOB W.: An adaptive parameterization for efficient material acquisition and rendering. *Transactions on Graphics (Proceedings of SIGGRAPH Asia) 37*, 6 (Nov. 2018), 274:1–274:18. doi:10.1145/3272127.3275059. 2, 4

[FV14] FILIP J., VÁVRA R.: Template-based sampling of anisotropic BRDFs. *Computer Graphics Forum 33*, 7 (October 2014), 91–99. doi:10.1111/cgf.12477. 2, 4

[Hei14] HEITZ E.: Understanding the masking-shadowing function in microfacet-based brdfs. *Journal of Computer Graphics Techniques (JCGT) 3*, 2 (June 2014), 48–107. 1

[LRR04a] LAWRENCE J., RUSINKIEWICZ S., RAMAMOORTHI R.: Efficient brdf importance sampling using a factored representation. *ACM Trans. Graph. 23*, 3 (Aug. 2004), 496–505. URL: https://doi.org/10.1145/1015706.1015751, doi:10.1145/1015706.1015751. 4

[LRR04b] LAWRENCE J., RUSINKIEWICZ S., RAMAMOORTHI R.: Efficient brdf importance sampling using a factored representation. In *ACM SIGGRAPH 2004 Papers* (New York, NY, USA, 2004), SIGGRAPH '04,

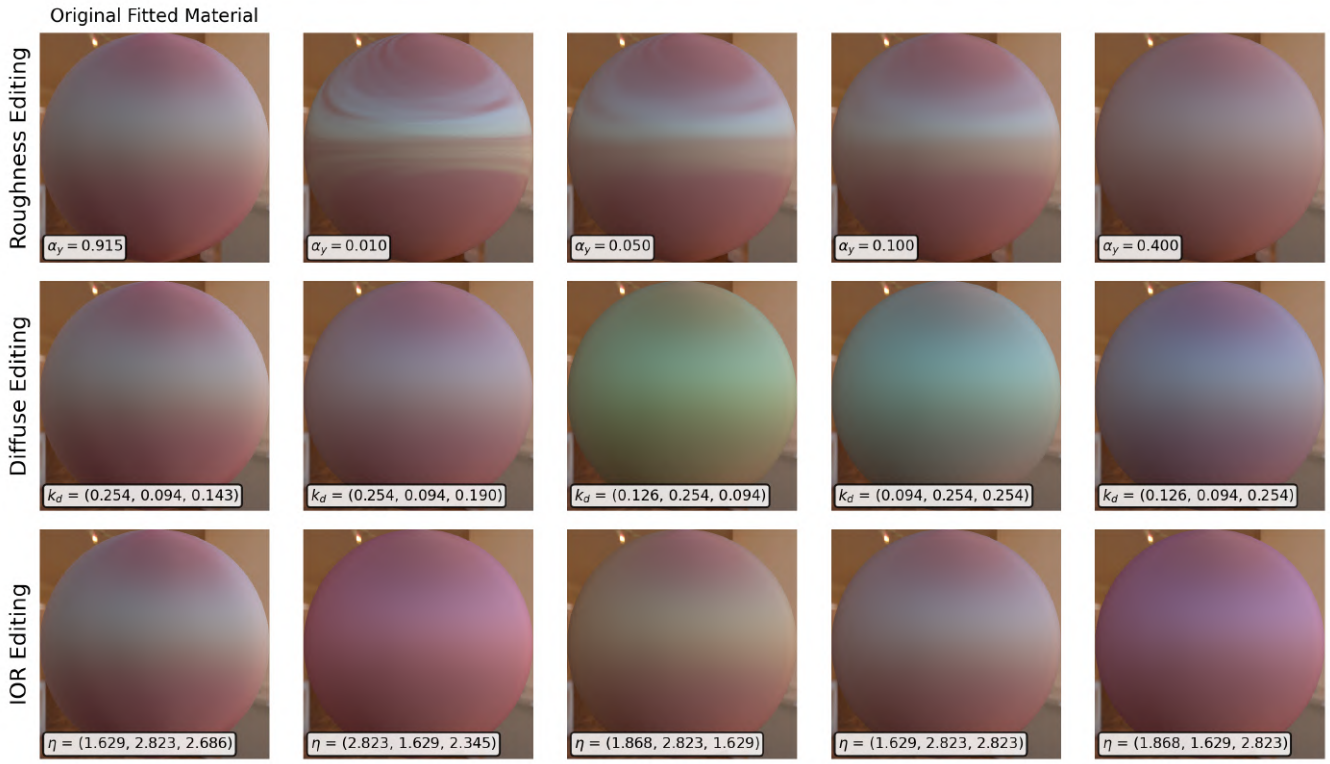

Figure 3: Editing an UTIA BRDF (`m071_fabric106`)

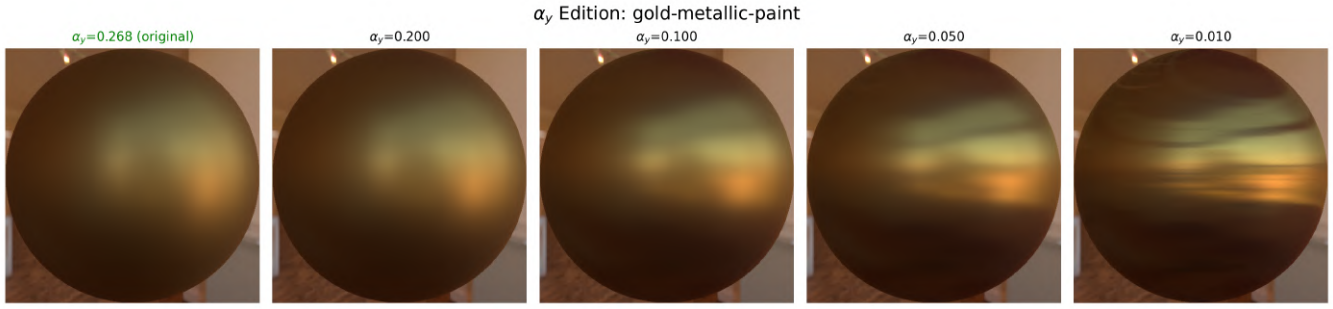

Figure 4: Making a MERL BRDF anisotropic after training.

Association for Computing Machinery, p. 496–505. doi:10.1145/1186562.1015751. 21, 22

[MPBM03] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: A data-driven reflectance model. In *ACM SIGGRAPH 2003 Papers* (New York, NY, USA, 2003), SIGGRAPH '03, Association for Computing Machinery, p. 759–769. URL: https://doi.org/10.1145/1201775.882343, doi:10.1145/1201775.882343. 1, 4

[Rus98] RUSINKIEWICZ S. M.: A new change of variables for efficient brdf representation. In *Eurographics Workshop on Rendering Techniques* (1998), Springer, pp. 11–22. 1

[VSW*23] VAIDYANATHAN K., SALVI M., WRONSKI B., AKENINE-MÖLLER T., EBELIN P., LEFOHN A.: Random-Access Neural Compression of Material Textures. In *Proceedings of SIGGRAPH* (2023). 1

[WMLT07] WALTER B., MARSCHNER S. R., LI H., TORRANCE K. E.: Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques* (Goslar, DEU, June 2007), EGSR'07, Eurographics Association, pp. 195–206. 1

[ZRW*24] ZELTNER T., ROUSSELLE F., WEIDLICH A., CLARBERG P., NOVÁK J., BITTERLI B., EVANS A., DAVIDOVIČ T., KALLWEIT S., LEFOHN A.: Real-time Neural Appearance Models. *ACM Transactions on Graphics 43*, 3 (June 2024), 1–17. doi:10.1145/3659577. 2

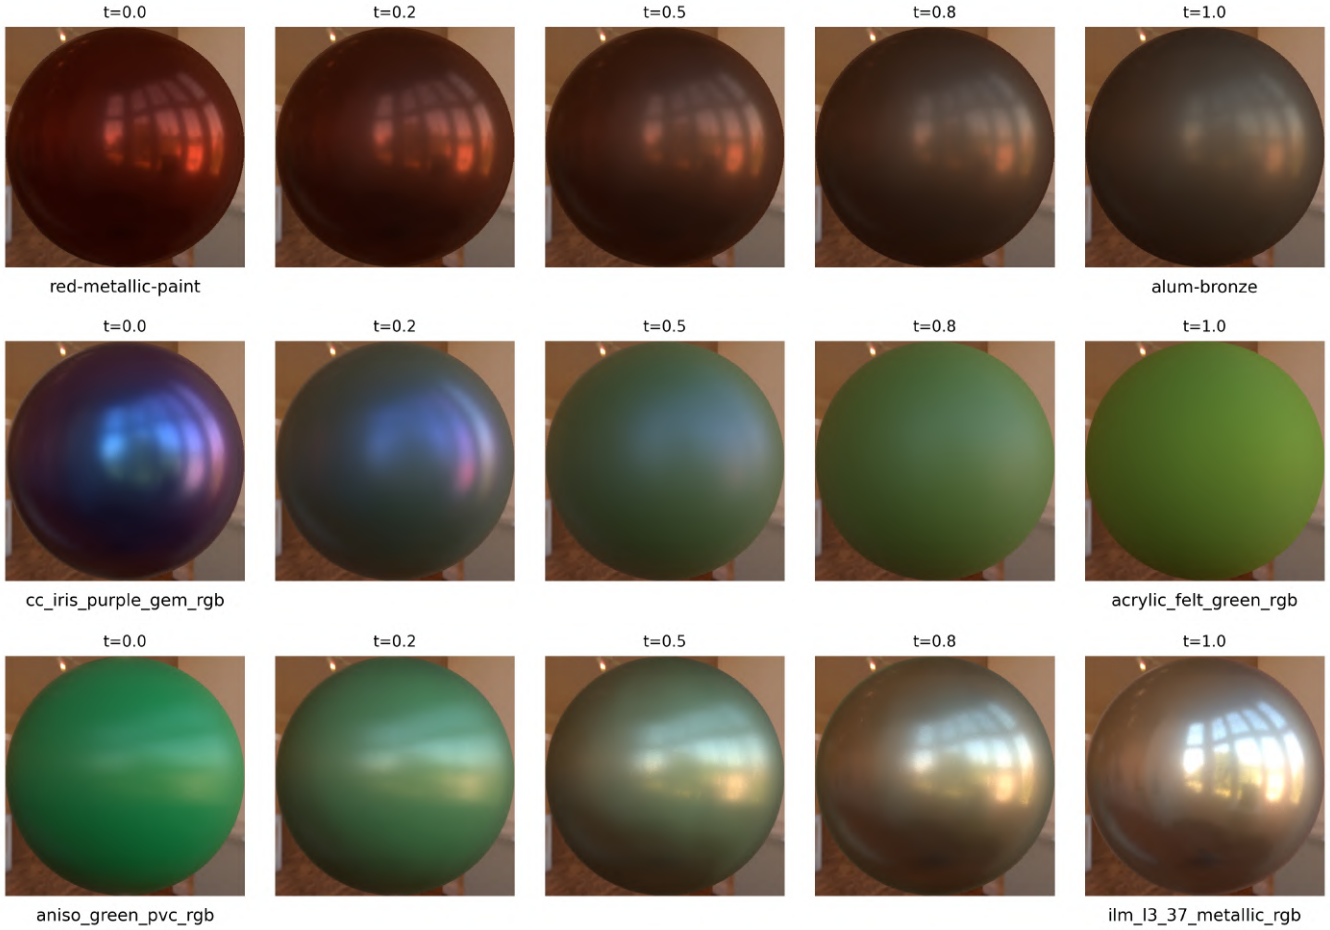

Figure 5: BRDF interpolation example. From left to right, we linearly interpolate (parameter t above) between two materials (MERL for the first line, and RGL for the second and third line), the latent vector, and the analytical parameters of our hybrid model, showing very plausible results.

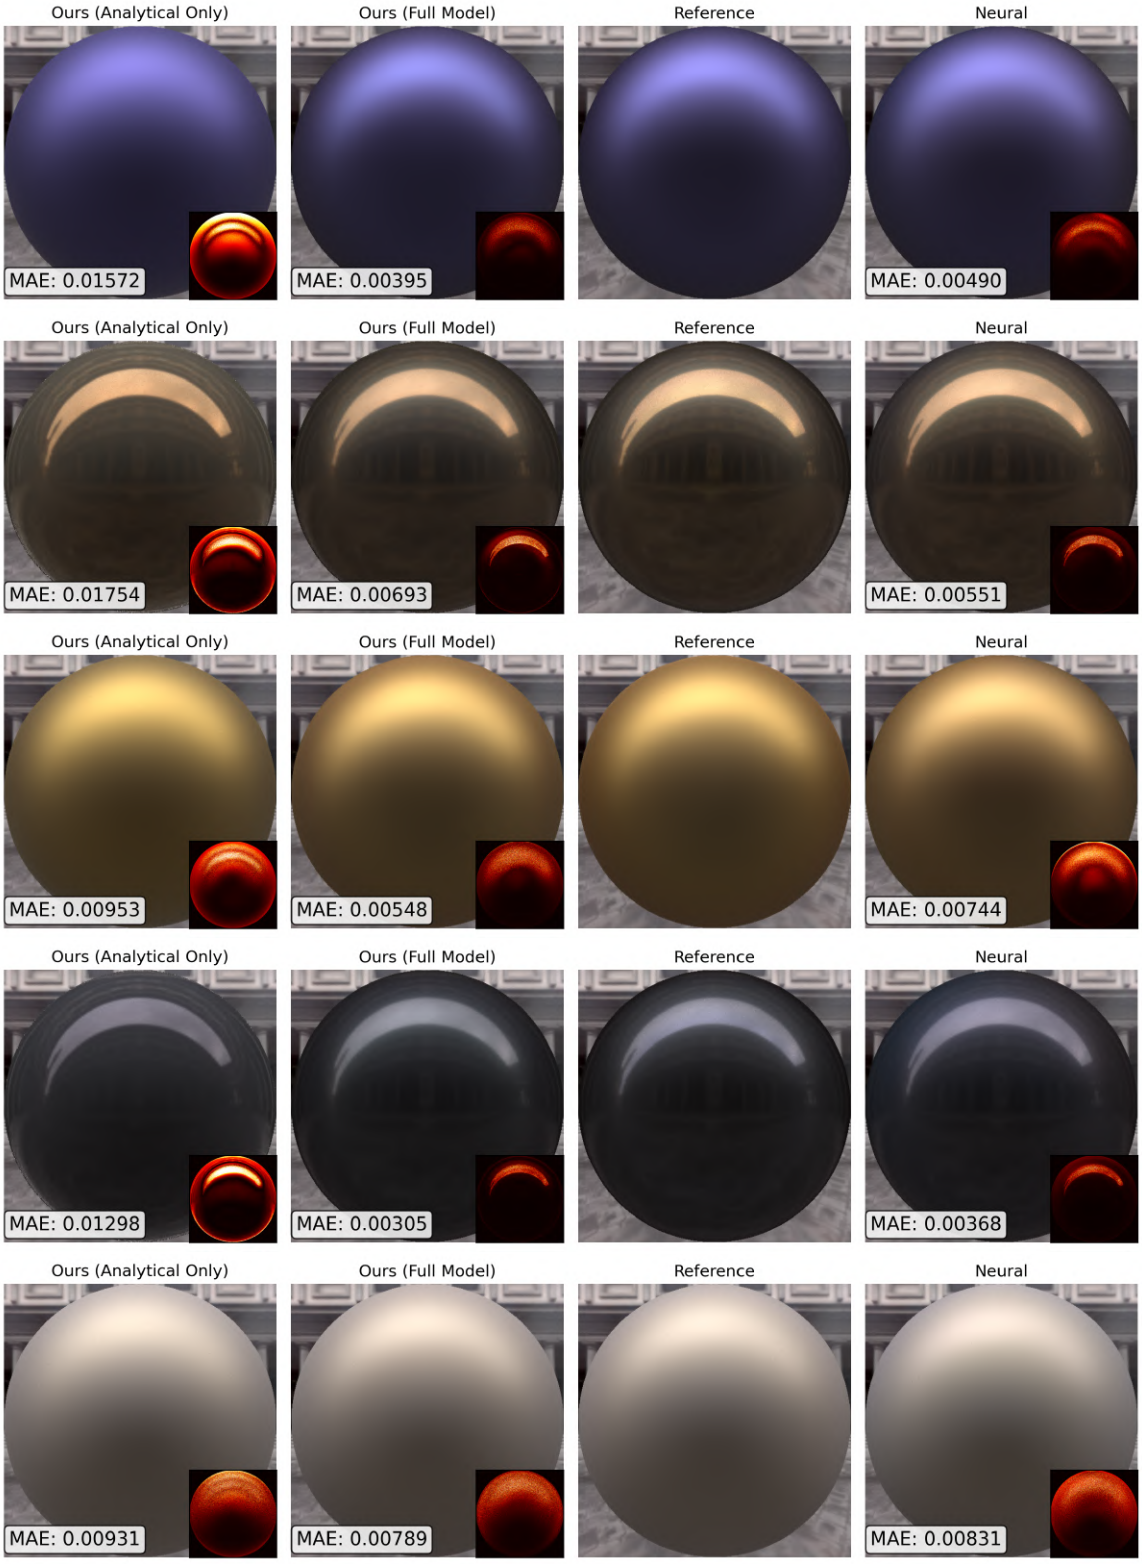

Figure 6: Results on the MERL dataset. **From top to bottom**: `blue-metallic-paint`, `brass`, `gold-metallic-paint`, `hematite` and `pearl-paint`. The MLP size is 32×3, and the latent vector is set to $l = 4$ for our hybrid model and $l = 12$ for the Neural model.

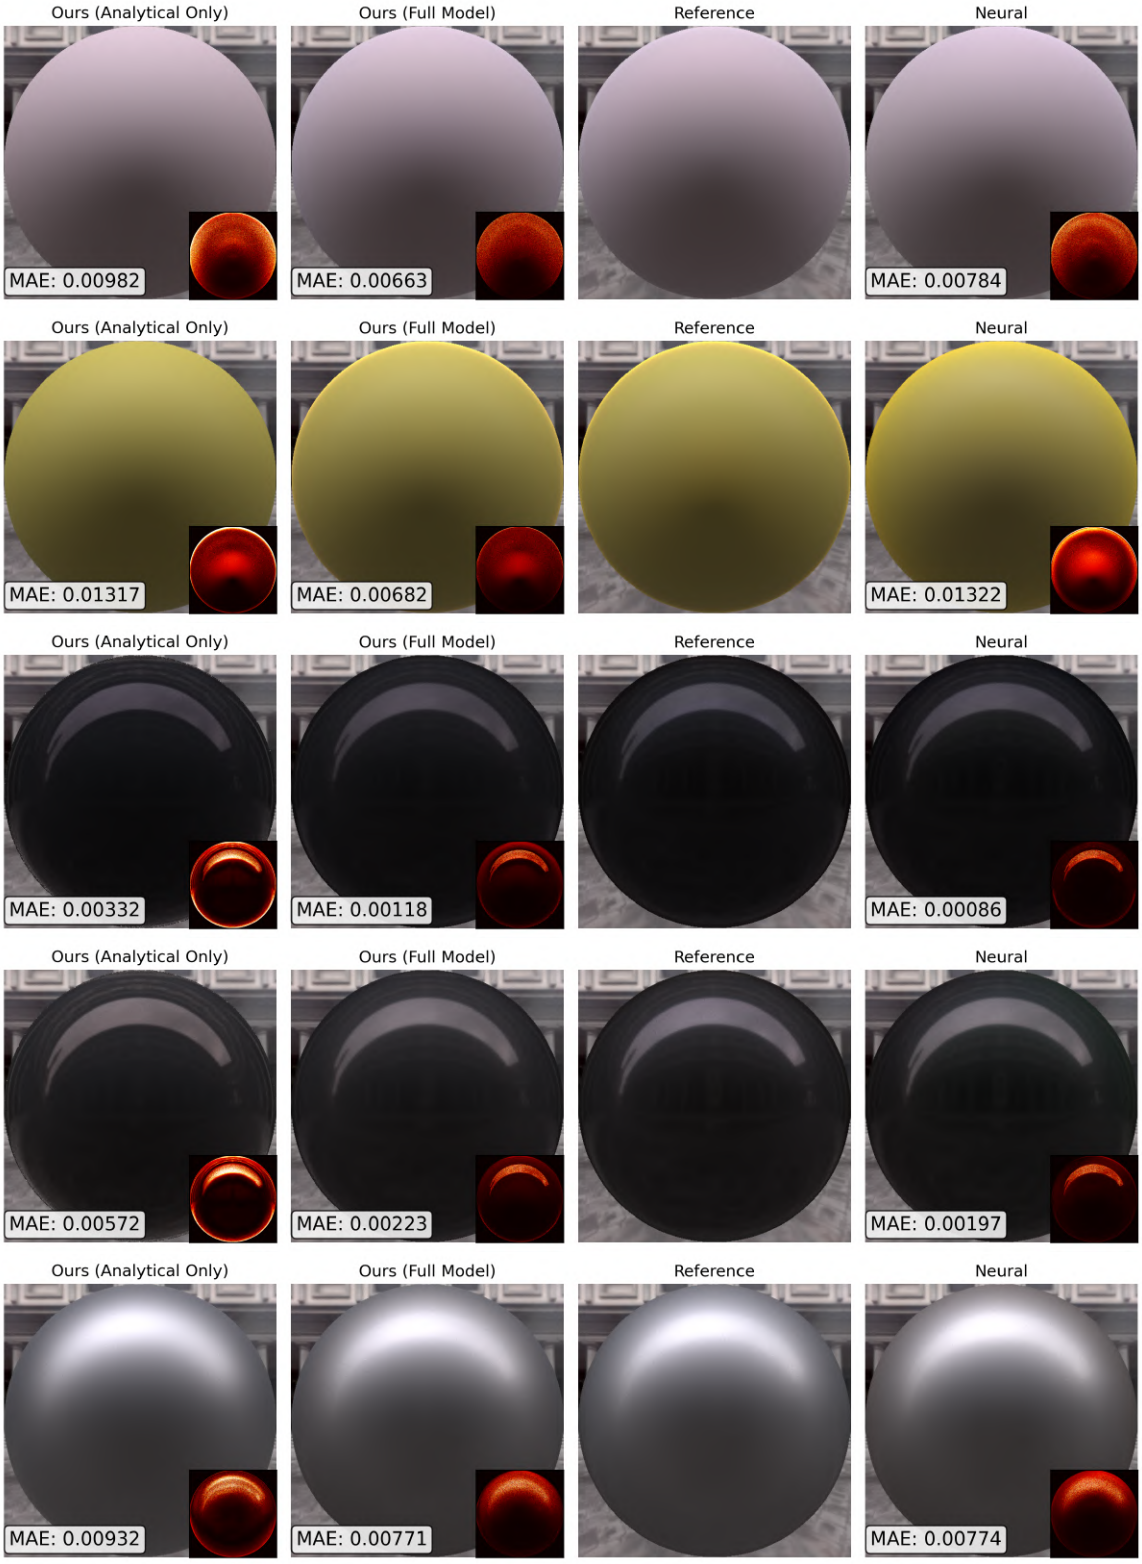

Figure 7: Other results on the MERL dataset. **From top to bottom**:`pink-fabric`, `yellow-plastic`, `specular-black-phenolic`, `silicon-nitrade` and `two-layer-silver`. The MLP size is 32×3, and the latent vector is set to $l = 4$ for our hybrid model and $l = 12$ for the Neural model.

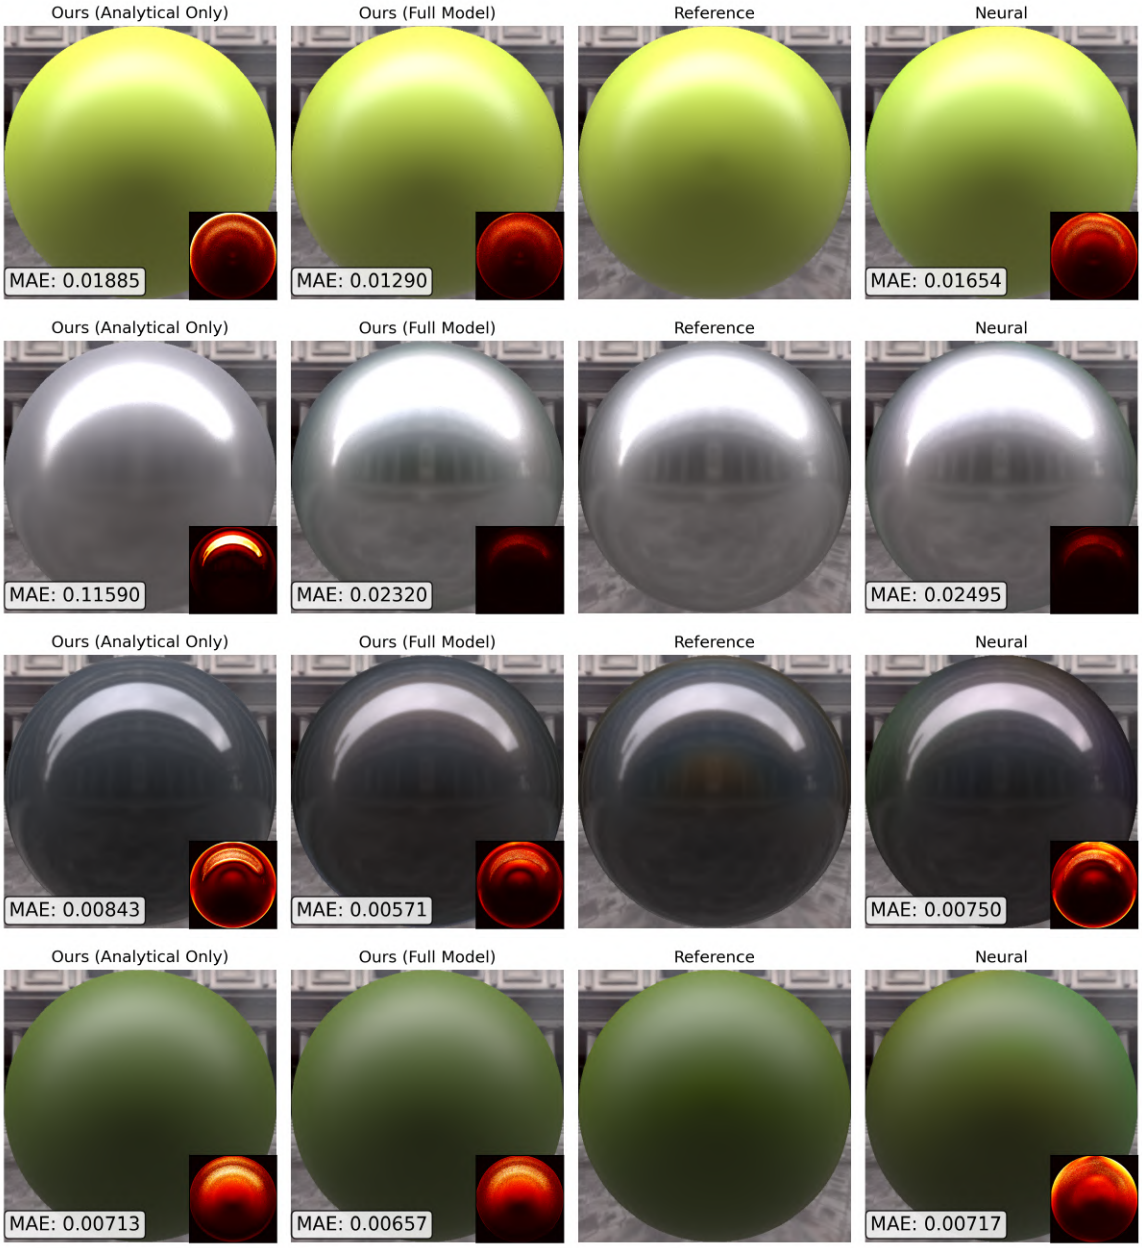

Figure 8: Results on the isotropic RGL dataset. **From top to bottom**: `cm_toxic_green`, `ilm_l3_37_metallic`, `irid_flake_paint1` and `leaf_maple`. The MLP size is 32×3, and the latent vector is set to $l = 4$ for our hybrid model and $l = 12$ for the Neural model.

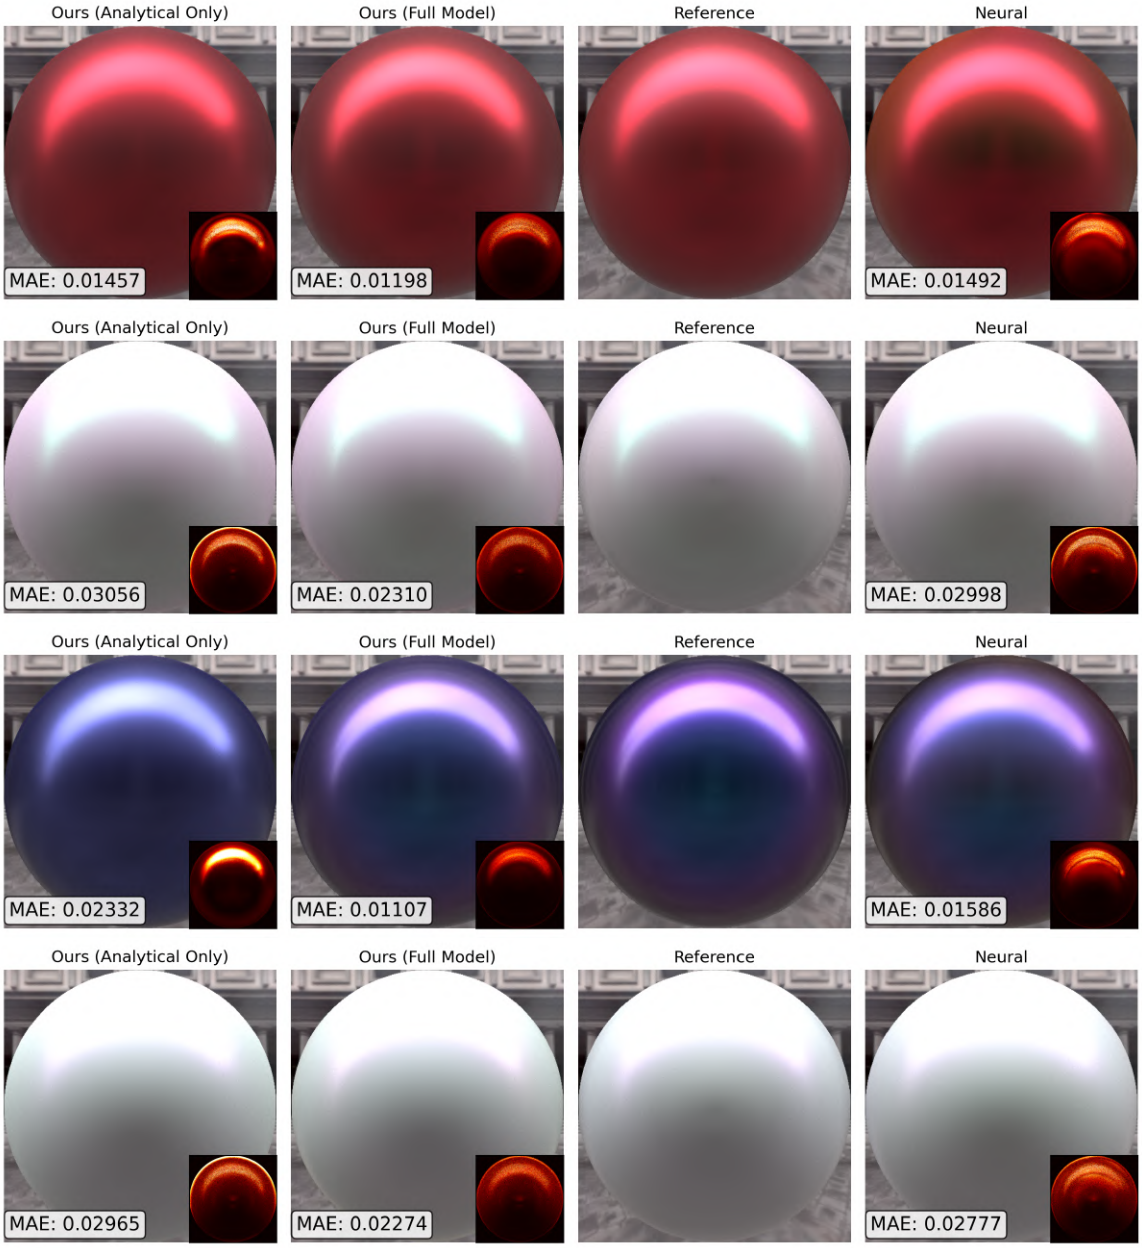

Figure 9: Other results on the isotropic RGL dataset. **From top to bottom**: `vch_dragon_eye_red`, `aurora_white`, `cc_iris_purple_gem` and `satin_purple`. The MLP size is 32×3, and the latent vector is set to $l = 4$ for our hybrid model and $l = 12$ for the Neural model.

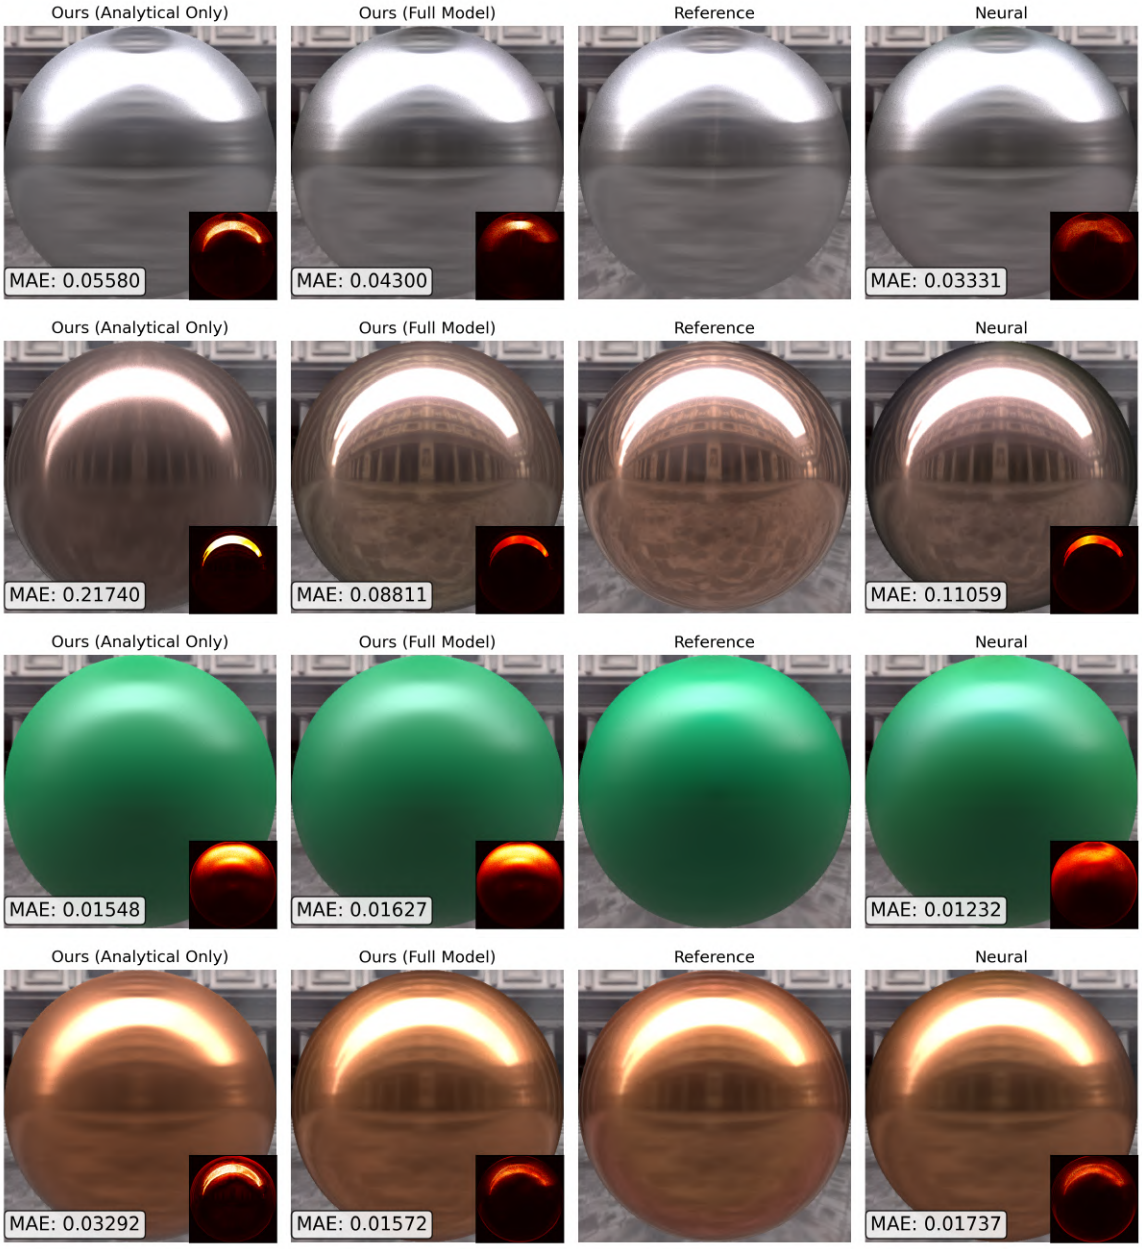

Figure 10: Results on the anisotropic RGL dataset.**From top to bottom**: `aniso_brushed_aluminium_1`, `aniso_copper_sheet`, `aniso_green_pvc` and `aniso_metallic_paper_copper`. The MLP size is 32×3, and the latent vector is set to $l = 4$ for our hybrid model and $l = 12$ for the Neural model.

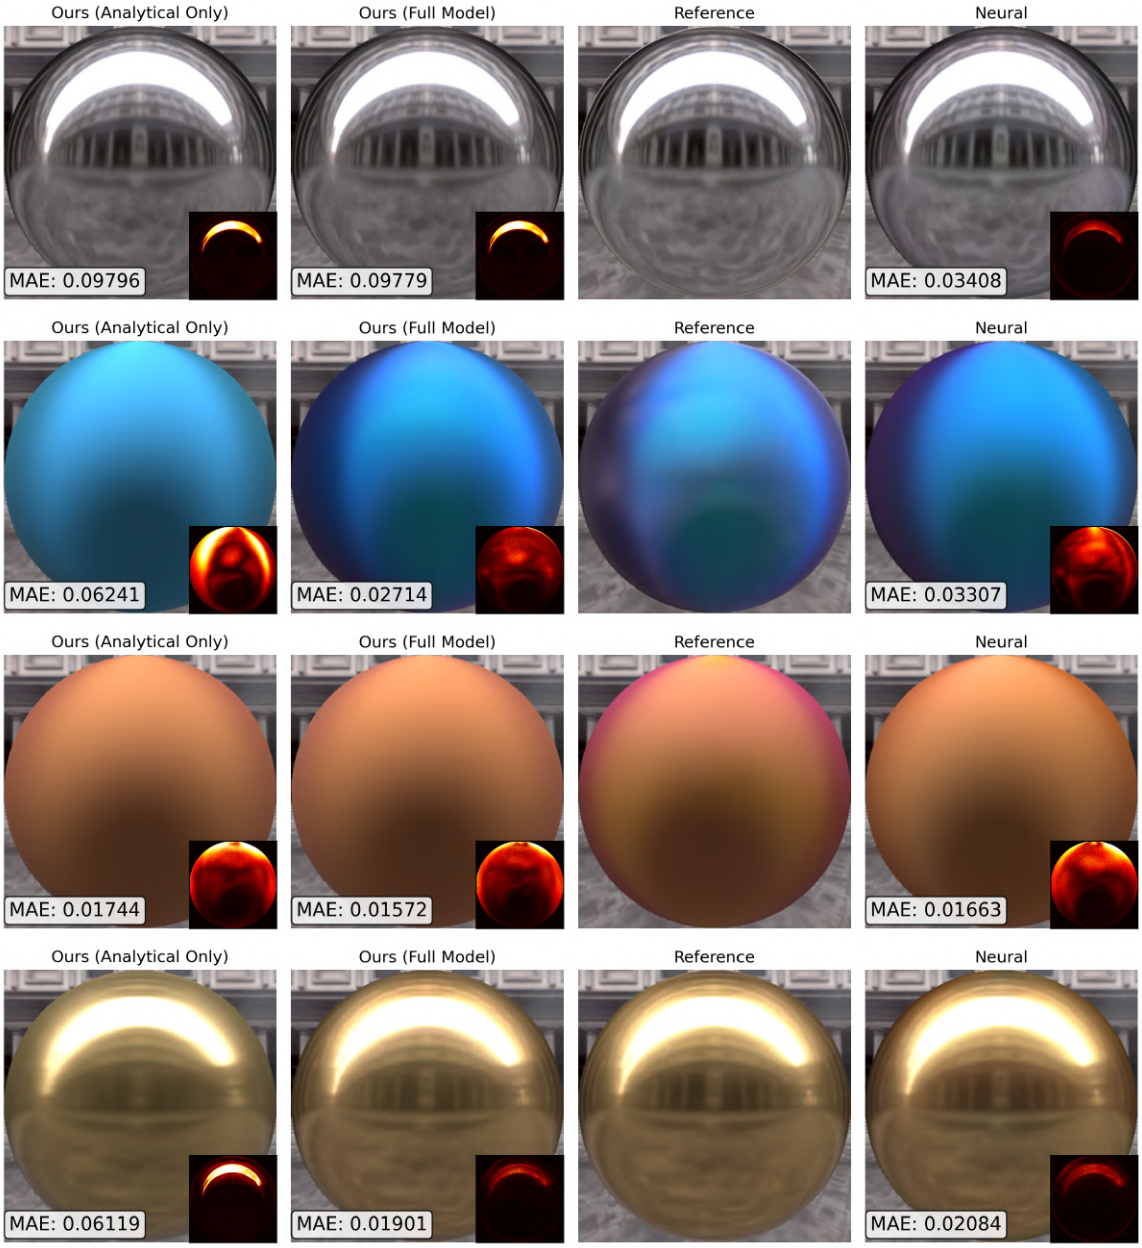

Figure 11: Other results on the anisotropic RGL dataset. **From top to bottom**: `aniso_miro_7`, `aniso_morpho_melenaus`, `aniso_sari_silk_2color` and `aniso_metallic_paper_gold`. The MLP size is 32×3, and the latent vector is set to $l = 4$ for our hybrid model and $l = 12$ for the Neural model.

Figure 12: Results on the UTIA dataset. `From top to bottom`: `m071_fabric106`, `m072_fabric107`, `m074_fabric111` and `m075_fabric112`. The MLP size is 32×3, and the latent vector is set to $l = 4$ for our hybrid model and $l = 12$ for the Neural model.

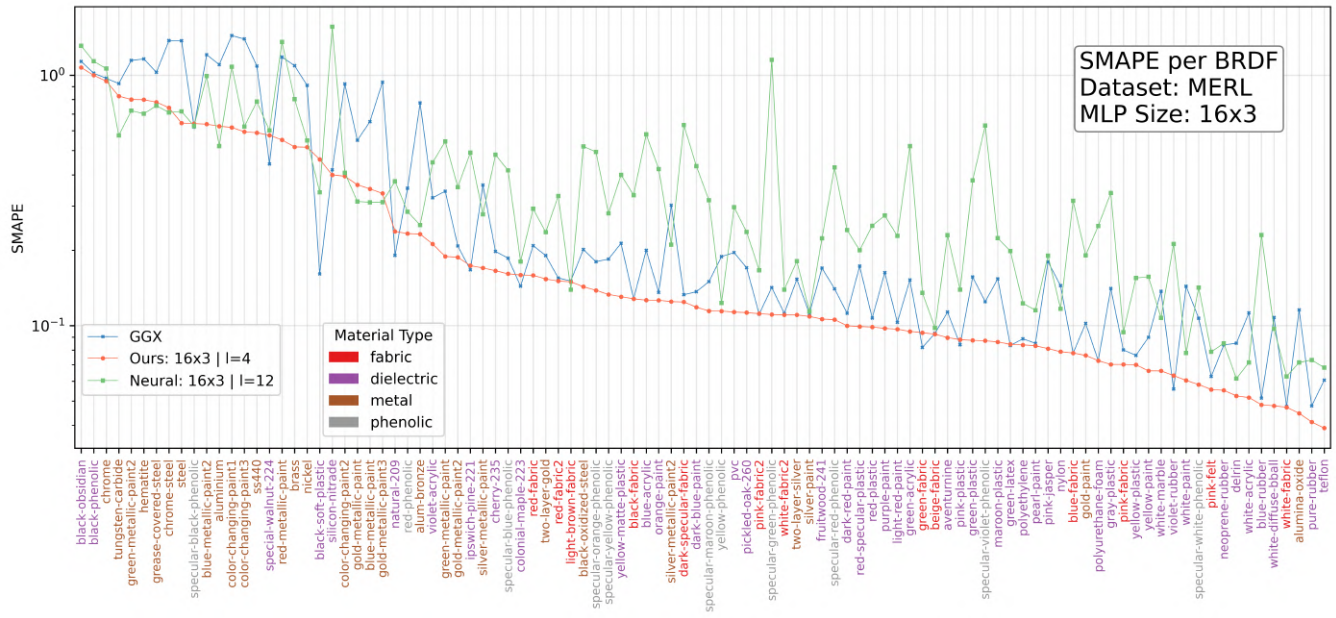

Figure 13: Comparisons, at constant memory cost, in BRDF space, in terms of SMAPE error, between the GGX analytical model, our hybrid model, and the Neural model. The MLP size is 16x3; our hybrid model and the neural one have the same number of parameters (12). Out of the 100 materials in the MERL database, our model outperforms other models in 88 of them. Note that for the materials `black-soft-plastic` and `natural-209`, the analytical GGX model outperforms our hybrid model as well as the neural one.

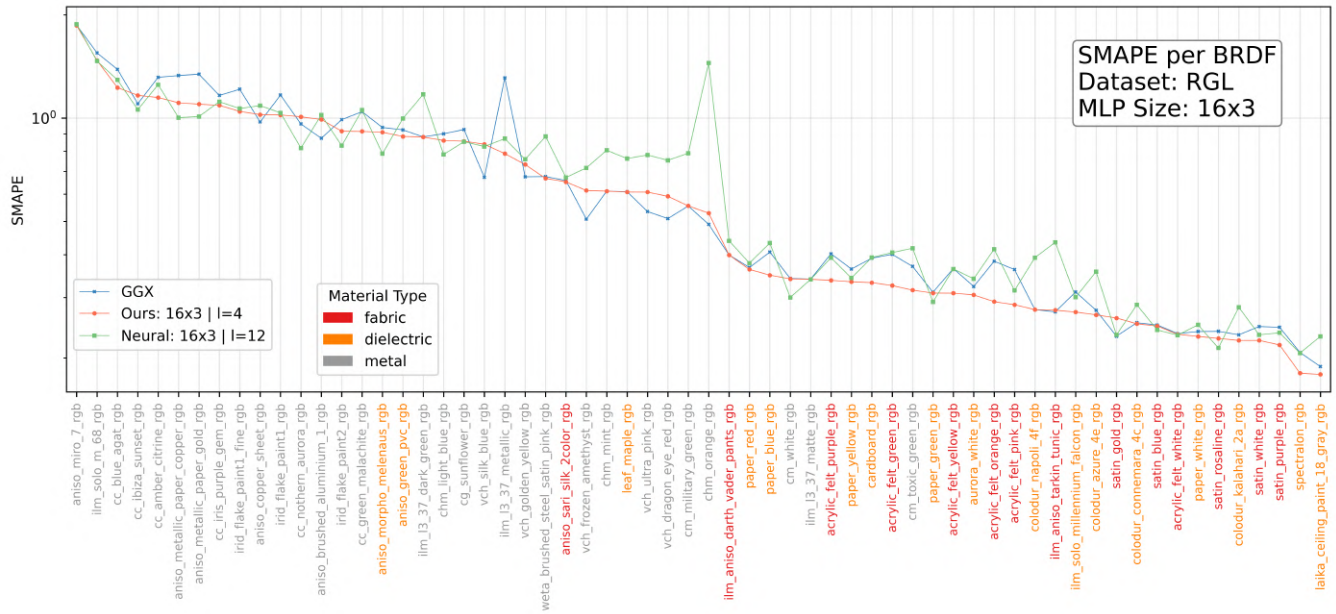

Figure 14: Comparisons, at constant memory cost, in BRDF space, in terms of SMAPE error, between the GGX analytical model, our hybrid model, and the Neural model. The MLP size is 16x3; our hybrid model and the neural one have the same number of parameters (12). Out of the 62 materials in the RGL dataset, our model outperforms other models in 50 of them.

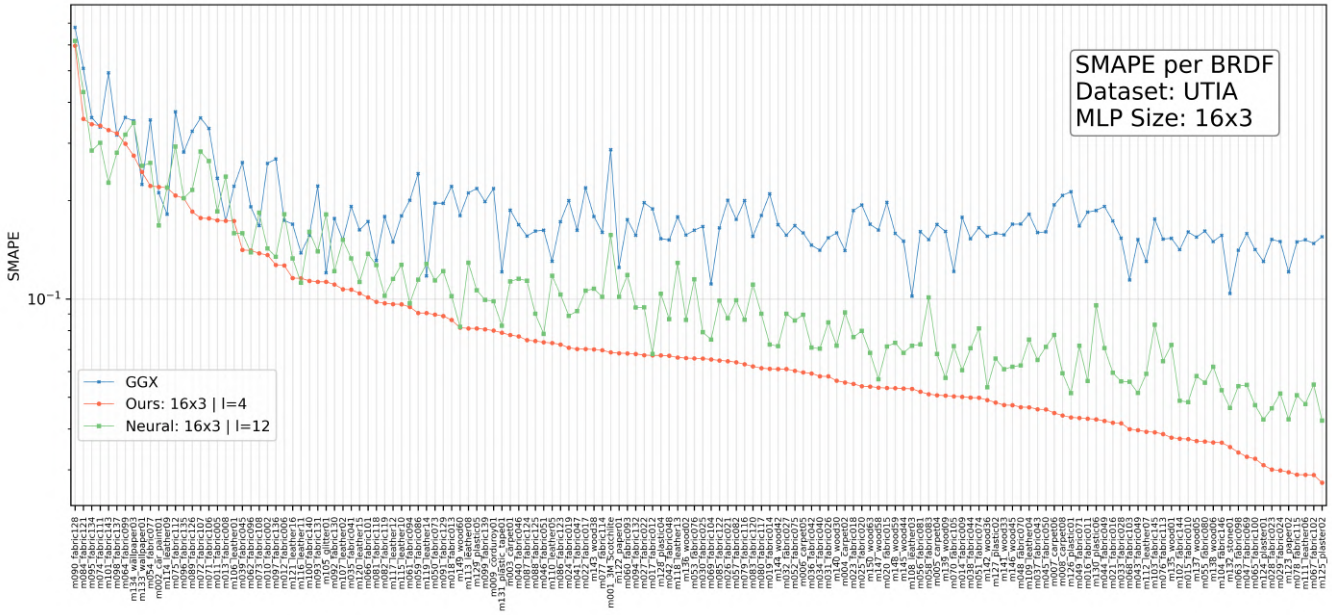

Figure 15: Comparisons, at constant memory cost, in BRDF space, in terms of SMAPE error, between the GGX analytical model, our hybrid model, and the Neural model. The MLP size is 16x3; our hybrid model and the neural one have the same number of parameters (13). Out of the 150 materials in the UTIA dataset, our model outperforms other models in 136 of them (only 14 materials are better fitted with the neural model).

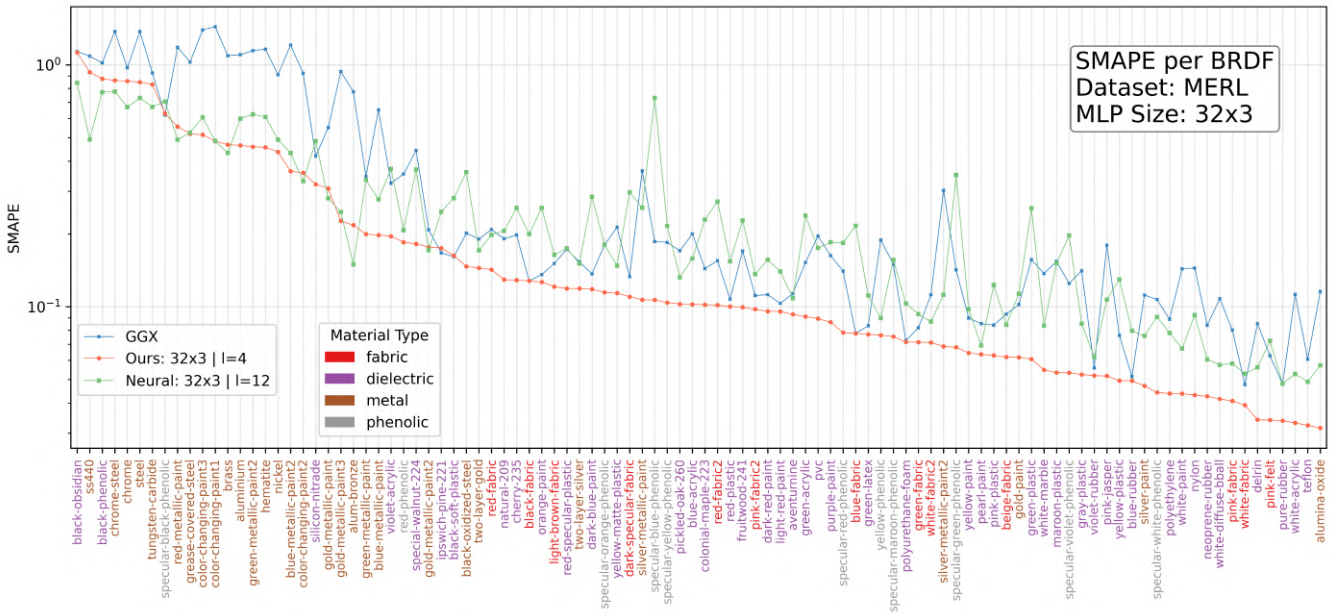

Figure 16: Comparisons, at constant memory cost, in BRDF space, in terms of SMAPE error, between the GGX analytical model, our hybrid model, and the Neural model. The MLP size is 32x3; our hybrid model and the neural one have the same number of parameters (12). Out of the 100 materials in the MERL database, our model outperforms other models in 88 of them.

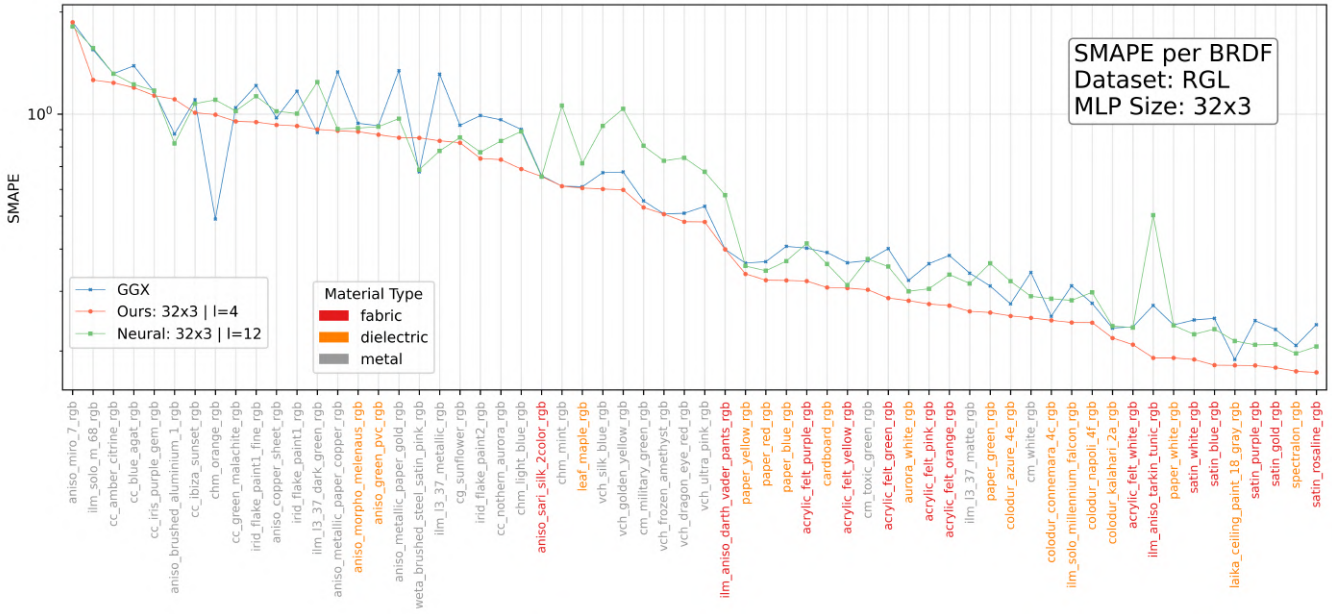

Figure 17: Comparisons, at constant memory cost, in BRDF space, in terms of SMAPE error, between the GGX analytical model, our hybrid model, and the Neural model. The MLP size is 32x3; our hybrid model and the neural one have the same number of parameters (12). Out of the 62 materials in the RGL dataset, our model outperforms other models in 58 of them.

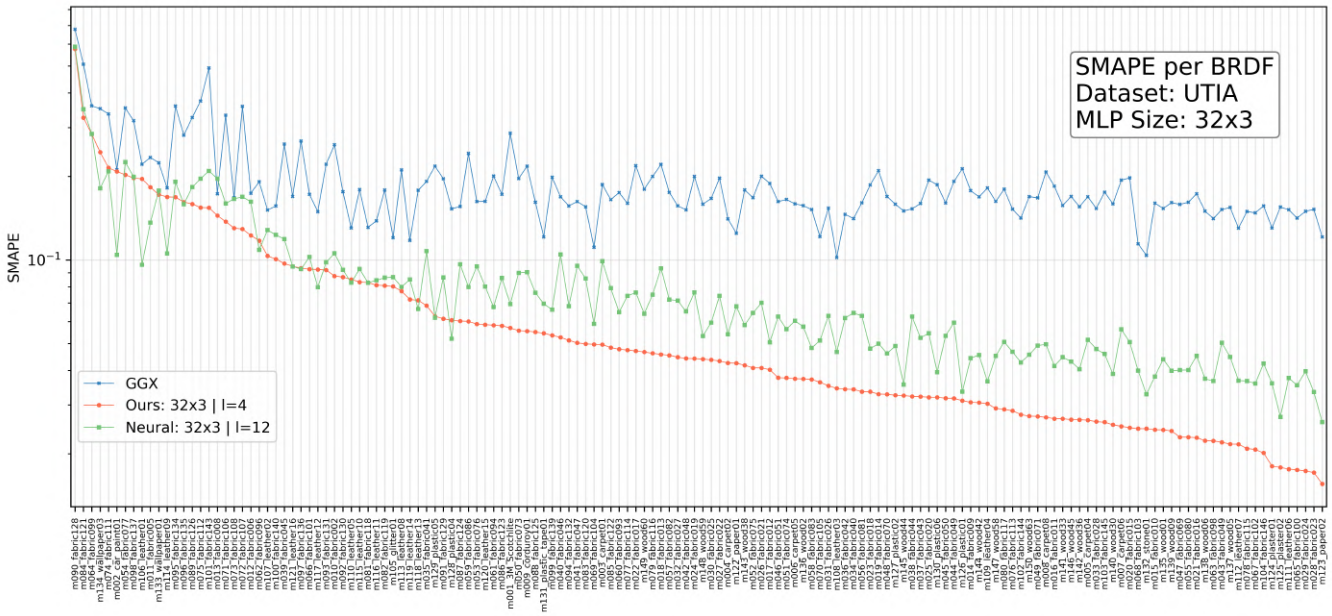

Figure 18: Comparisons, at constant memory cost, in BRDF space, in terms of SMAPE error, between the GGX analytical model, our hybrid model, and the Neural model. The MLP size is 32x3; our hybrid model and the neural one have the same number of parameters (13). Out of the 150 materials in the UTIA dataset, our model outperforms other models in 127 of them (only 23 materials are better fitted with the neural model).

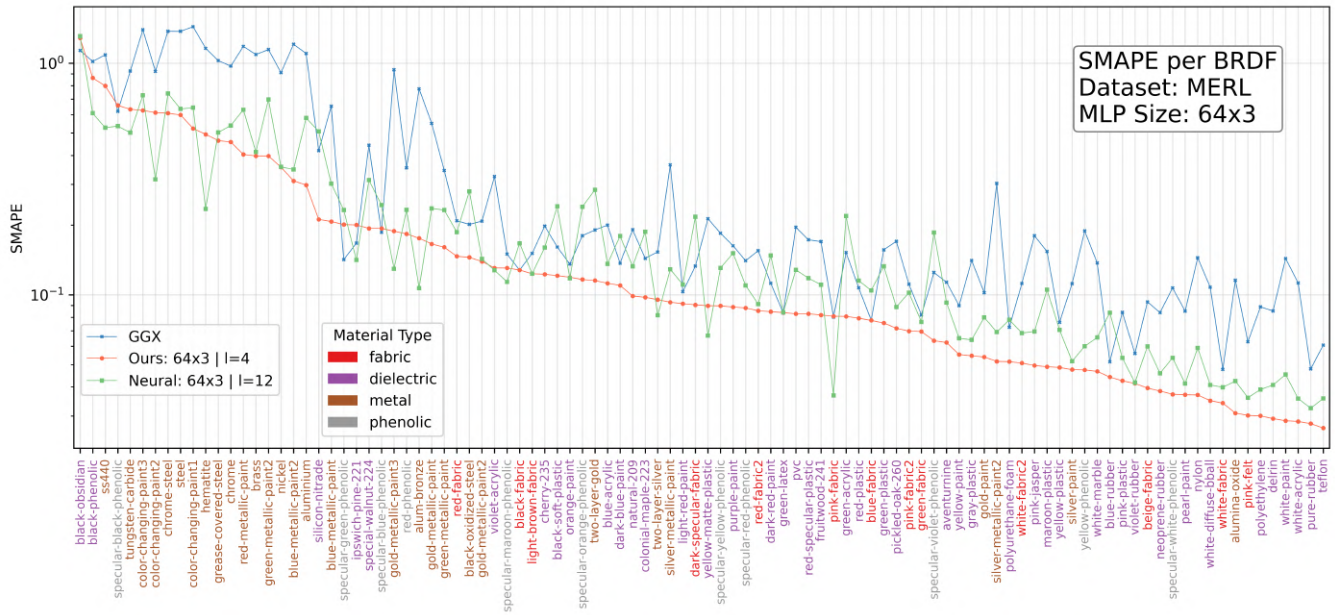

Figure 19: Comparisons, at constant memory cost, in BRDF space, in terms of SMAPE error, between the GGX analytical model, our hybrid model, and the Neural model. The MLP size is 64x3; our hybrid model and the neural one have the same number of parameters (12). Out of the 100 materials in the MERL database, our model outperforms other models in 85 of them.

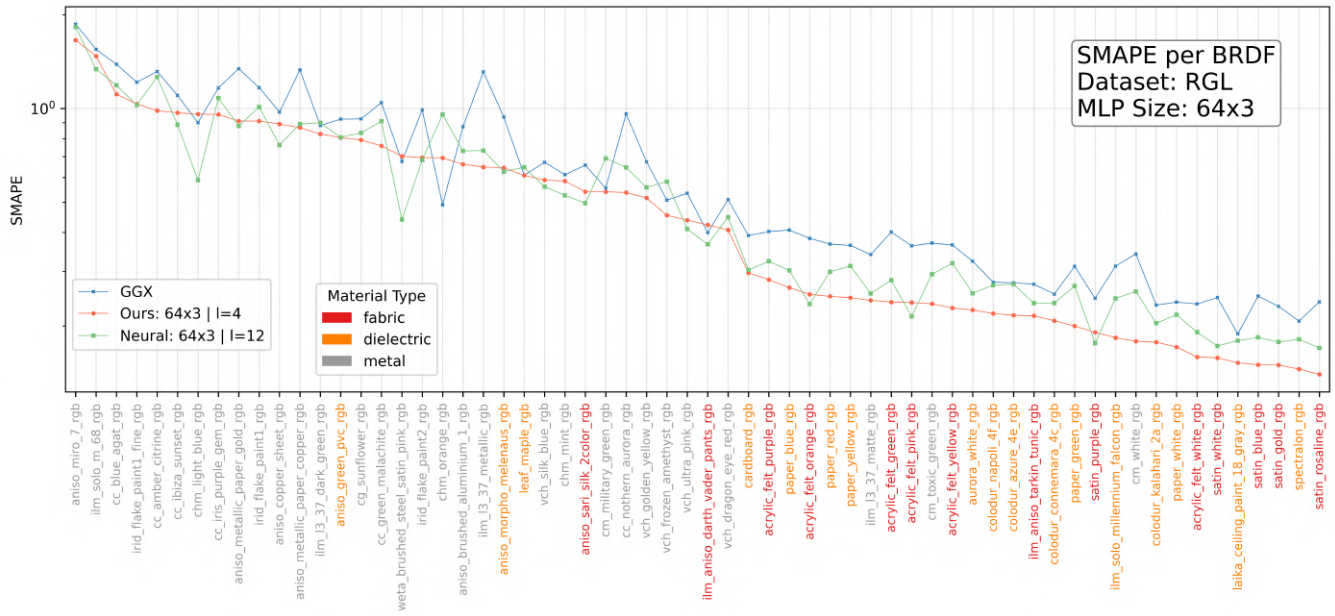

Figure 20: Comparisons, at constant memory cost, in BRDF space, in terms of SMAPE error, between the GGX analytical model, our hybrid model, and the Neural model. The MLP size is 64x3; our hybrid model and the neural one have the same number of parameters (12). Out of the 62 materials in the RGL dataset, our model outperforms other models in 43 of them.

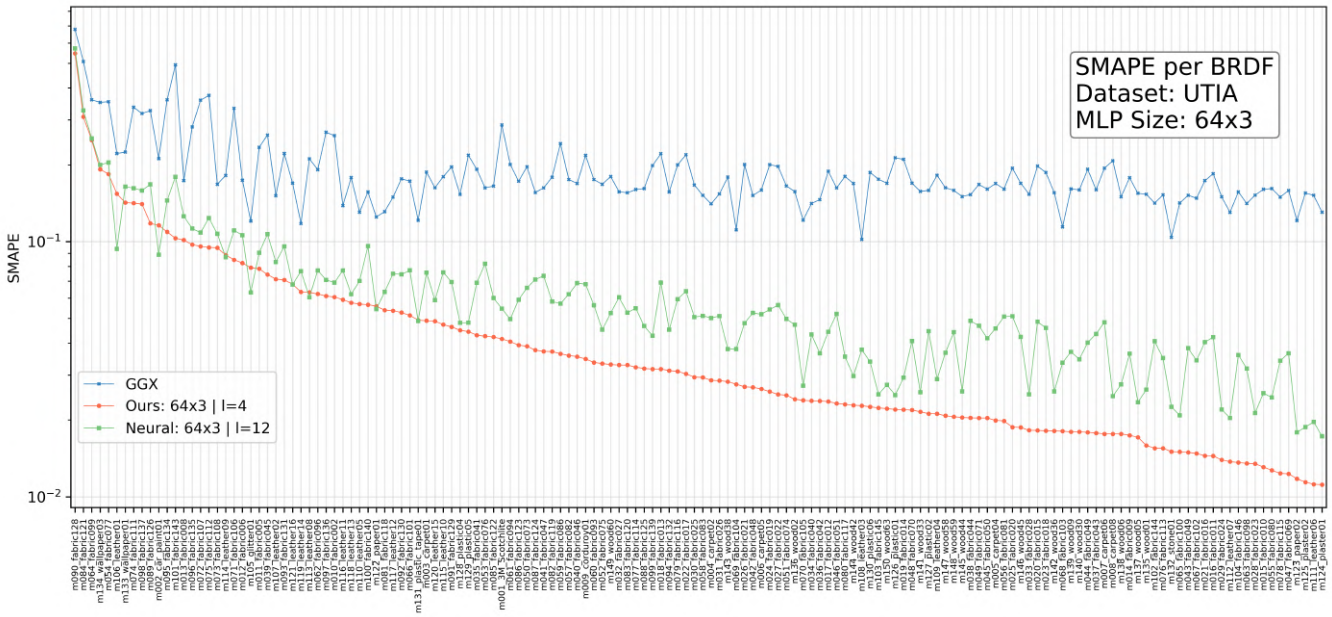

Figure 21: Comparisons, at constant memory cost, in BRDF space, in terms of SMAPE error, between the GGX analytical model, our hybrid model, and the Neural model. The MLP size is 64x3; our hybrid model and the neural one have the same number of parameters (13). Out of the 150 materials in the UTIA dataset, our model outperforms other models in 141 of them (only 9 materials are better fitted with the neural model).

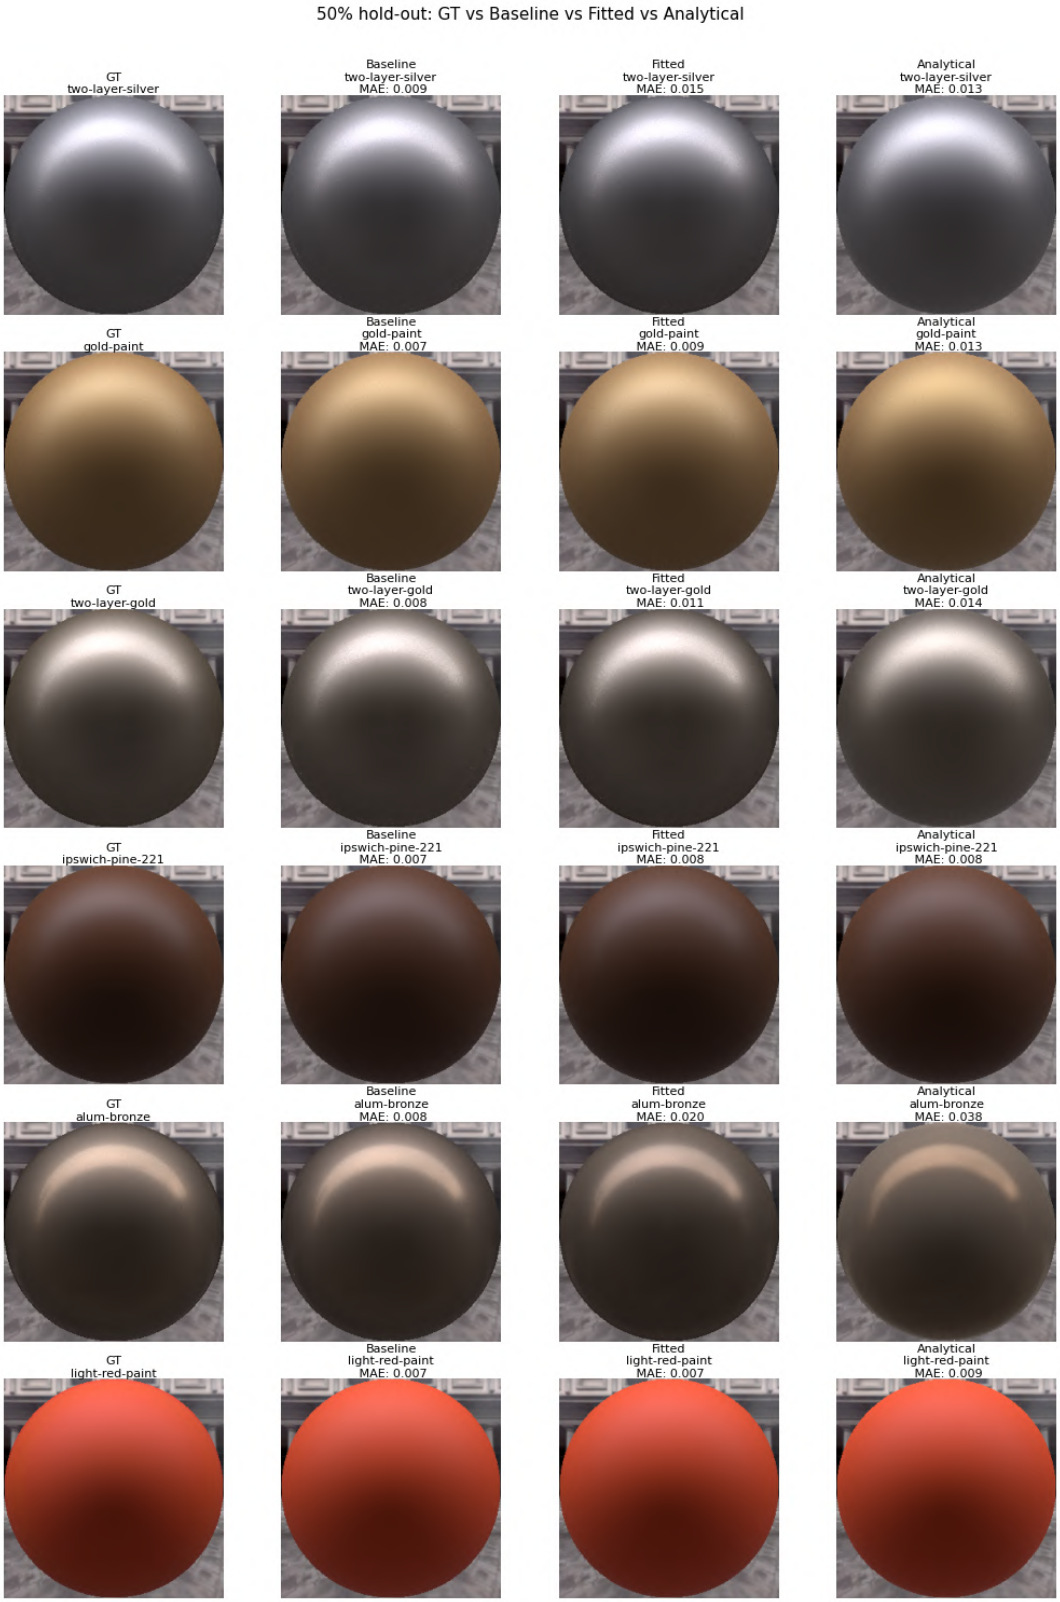

Figure 22: Rendering comparisons for BRDFs that were held out of the training set. We show the rendering MAE and compare them against the full training result and the analytical model.

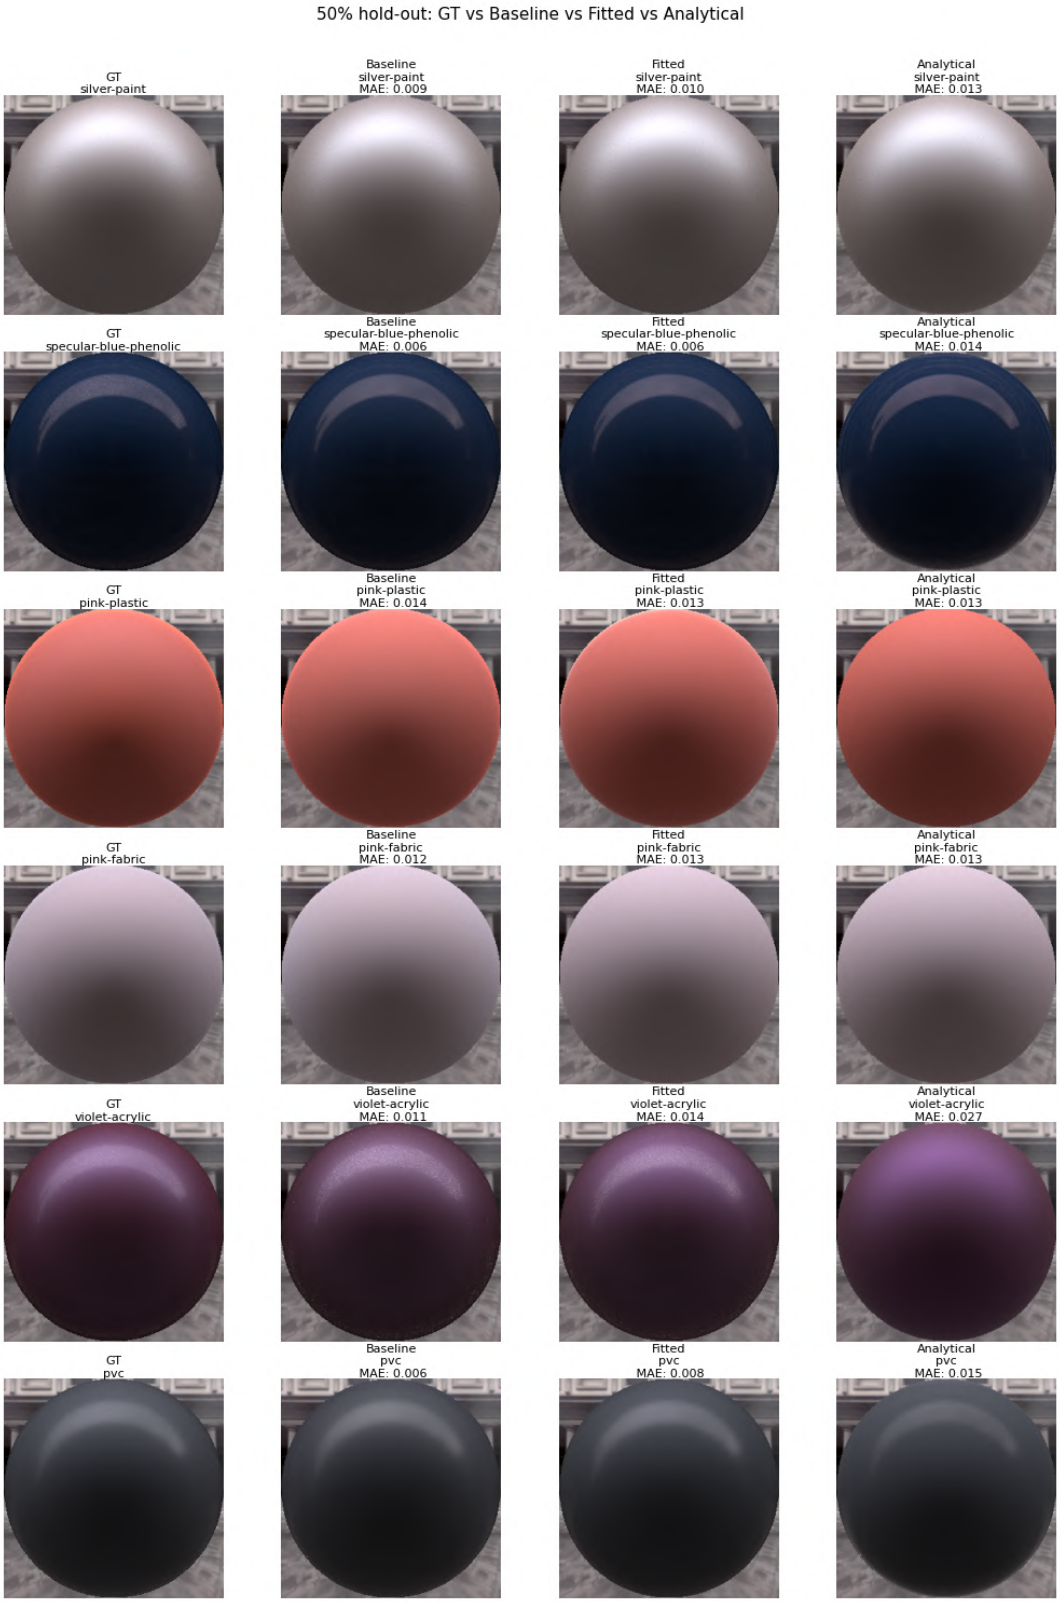

Figure 23: Rendering comparisons for BRDFs that were held out of the training set. We show the rendering MAE and compare them against the full training result and the analytical model.

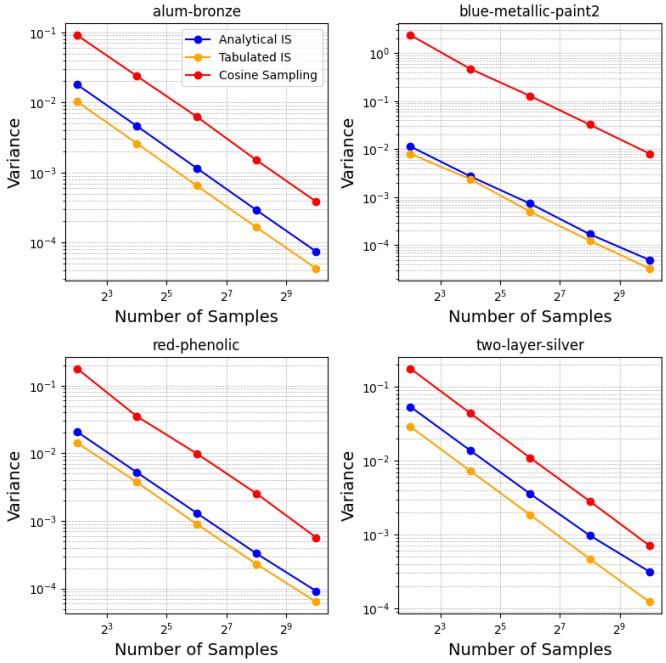

Figure 24: Color-averaged variance of the absolute difference between an image rendered with importance sampling and a reference image computed with 16k samples. Analytical Importance Sampling corresponds to applying either cosine sampling or GGX sampling using a MIS procedure for our model. As shown above, using our Analytical IS approaches the true PDF quite accurately compared to the reference solution provided by Tabulated Importance Sampling (cf. [LRR04b]).

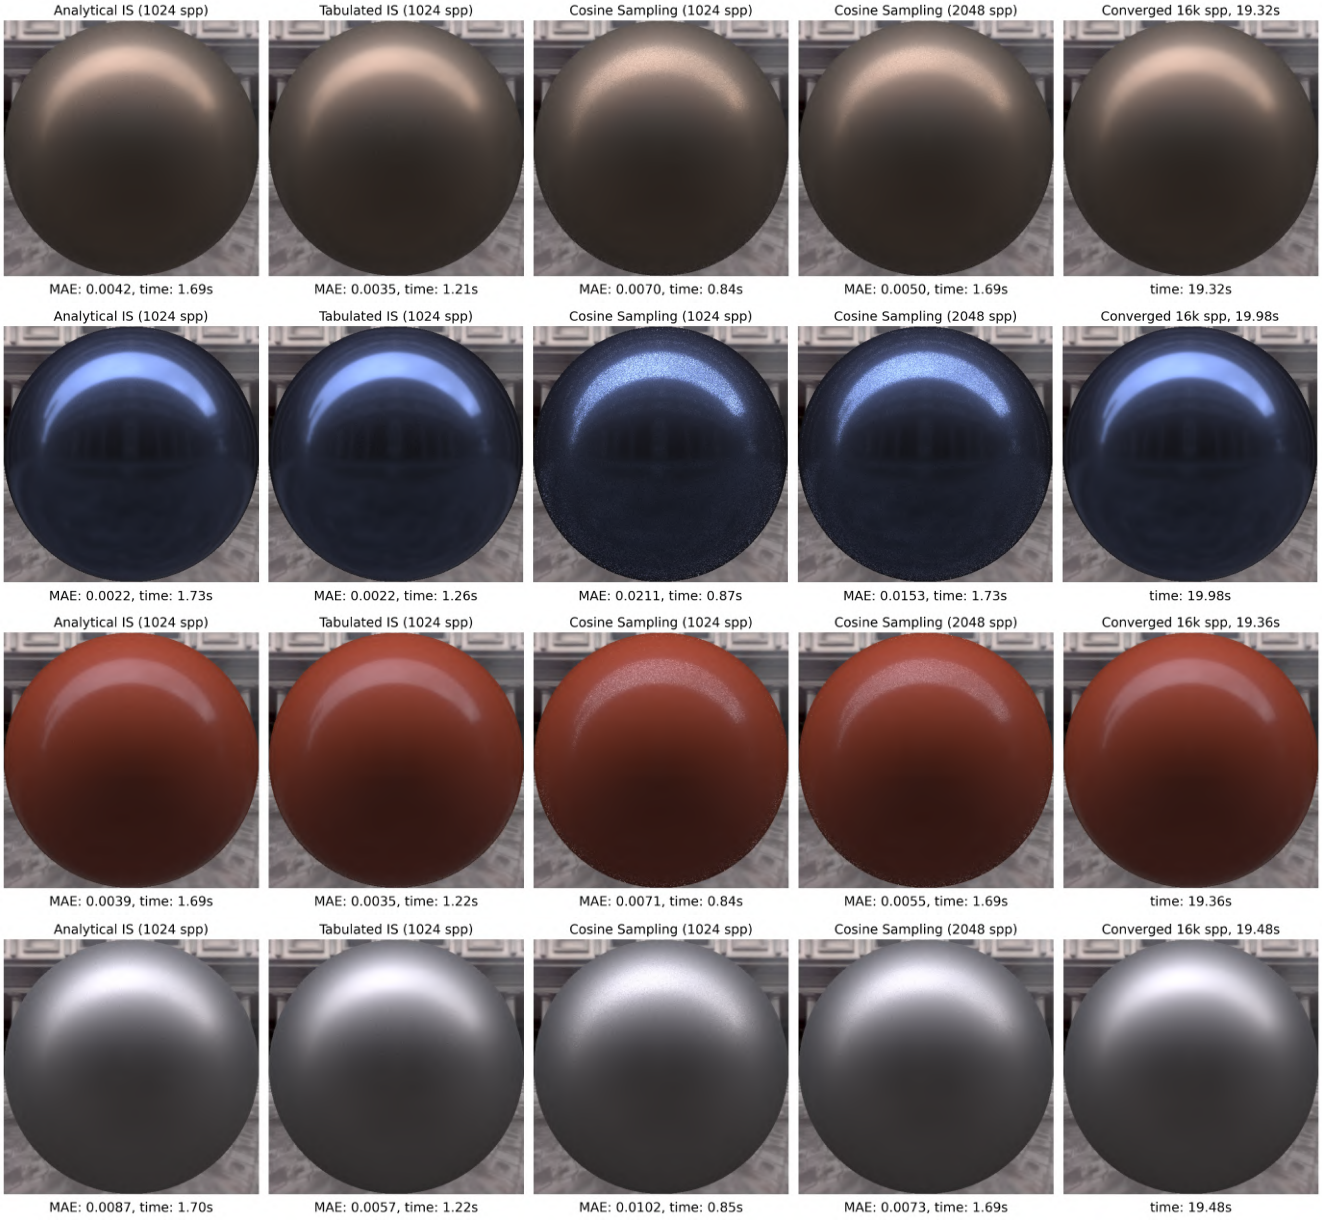

Figure 25: Importance Sampling (IS) comparisons. **From left to right**: IS based on the GGX distribution and a cosine lobe distribution, tabulated IS (similarly to [LRR04b]), with a memory footprint of 5.62MB), the cosine lobe IS, and a converged result using tabulated IS. For all images, the BRDF values are taken from our fitted hybrid model on the `alumn-bronze` material of the MERL dataset. As shown by the two images on the left, using an importance sampling based on the microfacet model is not as optimal as the tabulated form but provides clear visual improvements over cosine lobe importance sampling. **From top to bottom**: `alumn-bronze`, `blue-metallic-paint2`, `red-phenolic` and `two-layer-silver` materials from the MERL database.

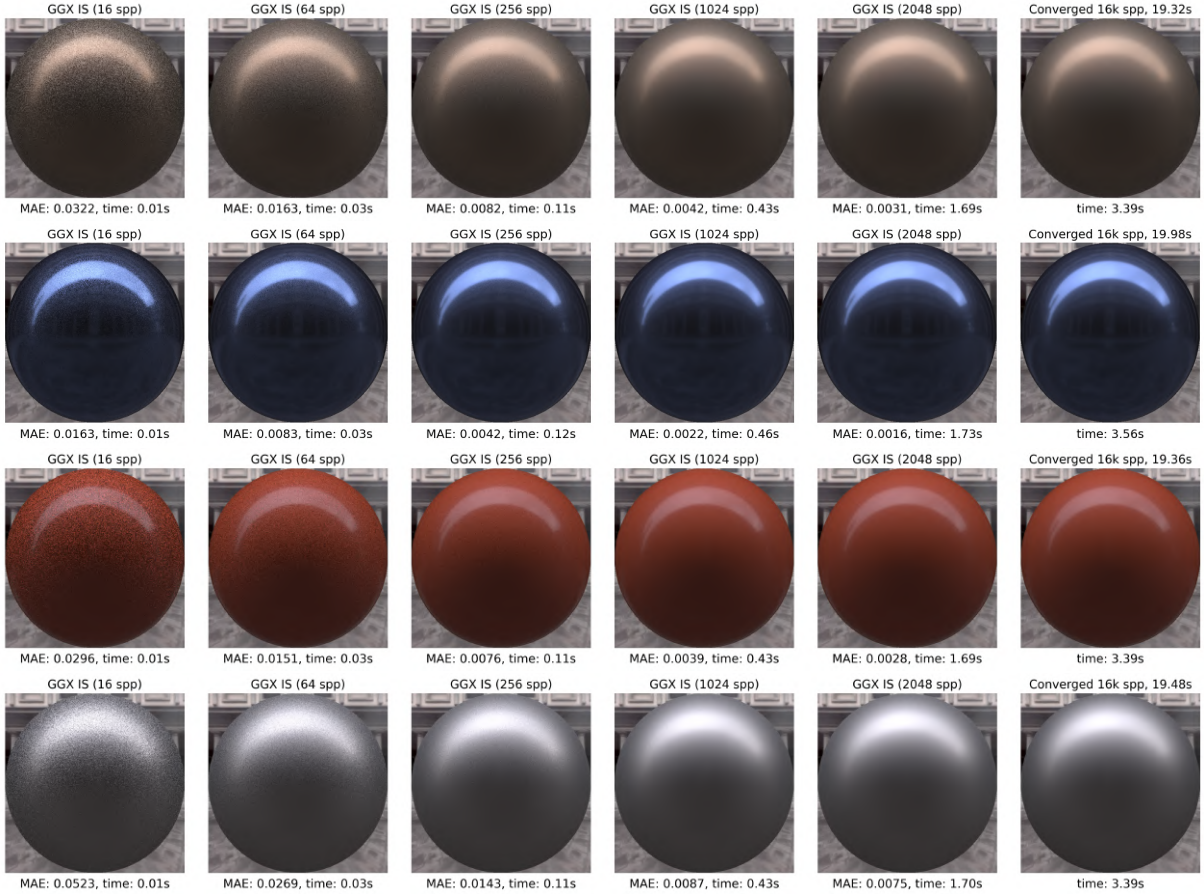

Figure 26: Rendered images obtained using an importance sampling method based only on the analytical component of our model (GGX distribution and a cosine lobe distribution) and corresponding to the data of Figure 24.

Figure 27: Path-tracing scene rendered using our Hybrid model for every visible object.