

Toward Richer Material Generation via Procedural Data Enhancement

Supplemental Document

Supplemental document for the paper, *Toward Richer Material Generation via Procedural Data Enhancement*. In this document, we discuss a list of topics that we did not expand on in the main paper, and also present a more comprehensive set of training data and generative results.

1 Non-Diffuse BRDF Lobes

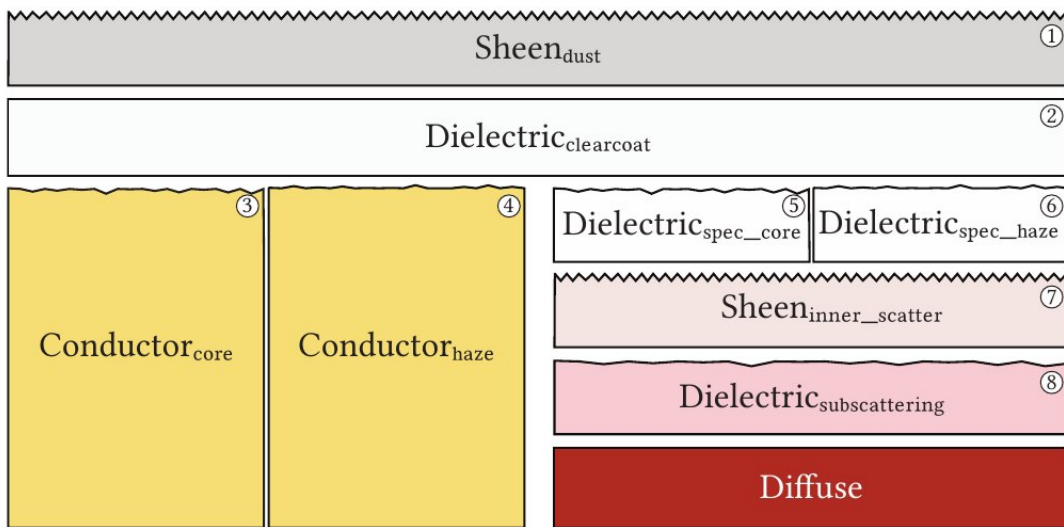


Figure 1: An illustration of our non-diffuse BSDF, consisting of eight lobes combined through mixing and coating. These lobes can be categorized into main specular reflection lobes and decorator lobes placed above and below the main material.

In this section, we describe our non-diffuse model and its components introduced in Section 3.1 of our paper, and include detailed design decisions. An illustration can be seen in Fig. 1. Here we outline the parameter ranges and their constraints, which are reflected in the training data of the universal MLP.

Our lobes can be roughly split into three different categories:

- The specular reflection lobes (dielectric/conductor core and haze), which form the main material
- Decorator lobes placed above the specular lobes (dust and clearcoat); these lobes affect the specular reflection lobes and may attenuate them
- Decorator lobes placed below the specular lobes (inner scatter and subcutaneous scatter); these lobes simulate low-order scattering from volumetric media below the surface of the main material and do not change the appearance of the main specular lobes

We layer these lobe components in an energy-conserving manner, where each successive lobe is modulated by the energy transmitted through the lobes above it, ensuring that lower layers cannot reflect more energy than they receive. This is similar to how this entire non-diffuse component is layered over the Lambertian diffuse term in our full BRDF, as described in the main paper.

Mathematically, our energy-conserving, multi-lobe BRDF model can be described as:

- Inner and subcutaneous scattering. This calculates the bottom-level BRDF from the dielectric branch of the material.

$$f_{\text{bottom}} = w_{\text{inner}} * f_{\text{sheen}}(t_{\text{inner}}, c_{\text{inner}}) + (1 - w_{\text{inner}} * A_{\text{inner}}) * w_{\text{subc}} * f_{\text{GGX}}(\eta_{\text{subc}}, \alpha_{\text{subc}}, c_{\text{subc}}) \quad (\text{S.1})$$

- Dielectric branch. We calculate the blended albedo (T_{diel}) used to transmit light down to the bottom layer, and then calculate the dielectric core and haze blend, and add the contribution from f_{bottom} .

$$\begin{aligned} T_{\text{diel}} &= (1 - m_{\text{ch}_d}) * (1 - A_{\text{d}_\text{core}}) + m_{\text{ch}_d} * (1 - A_{\text{d}_\text{haze}}) \\ f_{\text{diel}} &= m_{\text{ch}_d} * f_{\text{GGX}}(\eta_{\text{d}_\text{core}}, \alpha_{\text{d}_\text{core}}) + (1 - m_{\text{ch}_d}) * f_{\text{GGX}}(\eta_{\text{d}_\text{haze}}, \alpha_{\text{d}_\text{haze}}) + T_{\text{diel}} * f_{\text{bottom}} \end{aligned} \quad (\text{S.2})$$

- Conductor branch. This calculates the specular reflection from the conductive GGX lobes.

$$f_{\text{cond}} = m_{\text{ch}_c} * f_{\text{GGX}}(\eta_{\text{r}_\text{core}}, \eta_{\text{i}_\text{core}}, \alpha_{\text{c}_\text{core}}) + (1 - m_{\text{ch}_c}) * f_{\text{GGX}}(\eta_{\text{r}_\text{haze}}, \eta_{\text{i}_\text{haze}}, \alpha_{\text{c}_\text{haze}}) \quad (\text{S.3})$$

- Base layer. We now mix together the conductor and dielectric branches, forming the base material without the top decorator effects.

$$f_{\text{base}} = m_{\text{dc}} * f_{\text{cond}} + (1 - m_{\text{dc}}) * f_{\text{diel}} \quad (\text{S.4})$$

- Clearcoat layer. This evaluates the clearcoat lobe and transmits the remaining energy down to the base layer.

$$f_{\text{clear}} = w_{\text{clear}} * f_{\text{GGX}}(\eta_{\text{clear}}, \alpha_{\text{clear}}) + (1 - w_{\text{clear}} * A_{\text{clear}}) * f_{\text{base}} \quad (\text{S.5})$$

- Final BSDF (adding dust layer). Finally, we evaluate the top-most dust layer and transmit the remaining energy to the layers underneath.

$$f = w_{\text{dust}} * f_{\text{sheen}}(t_{\text{dust}}, c_{\text{dust}}) + (1 - w_{\text{dust}} * A_{\text{dust}}) * f_{\text{clear}} \quad (\text{S.6})$$

In the equations above, f denotes the respective BSDFs and A is the sum of the reflection albedo and the absorption of a BSDF, so that the transmission albedo A_t can be computed as $A_t = 1 - A$. Descriptions of the corresponding parameters are given in the following sections and in Table 4. For additional details, we refer the readers to the cited work.

Dielectric Core Lobe (ID: 5)

The dielectric core lobe is a GGX lobe [Walter et al., 2007] with a relatively low roughness α_{core} . The index of refraction (IOR) η_d drives the intensity of the lobe, and we determine the range of this IOR based on those of real, dielectric materials.

We decided to use the full Fresnel coefficient [Born and Wolf, 2013] for reflection rather than an empirical version commonly used in real-time rendering, because most dielectric materials have a low reflectance with a specular albedo of < 0.06 at normal incidence, and the classic Schlick approximation [Schlick, 1994] becomes inaccurate for IORs < 1.4 .

To summarize, the dielectric core lobe is described by the following parameters:

- Dielectric index of refraction, or IOR (η_d)
- Lobe roughness (α_{core})

Dielectric Haze Lobe (ID: 6)

The dielectric haze lobe is another GGX lobe and a “companion” lobe to the core lobe. Core and haze lobes share the same IOR because we want to simulate more complex microfacet-like distributions, rather than having two independent media that are mixed together.

The dielectric haze lobe is described by the following parameters:

- Dielectric IOR (η_d), identical to that from the core lobe
- Lobe roughness (α_{haze}), larger than the core roughness

Conductive Core Lobe (ID: 3)

This is a GGX lobe that uses a complex-valued IOR for its reflectance. We use the full, unpolarized Fresnel coefficients. The real and imaginary components of the IOR (η_r, η_i) are derived from an RGB color following [Gulbrandsen, 2014]. We choose to explicitly specify the reflection color (metal color) instead of using tabulated IOR values, because it makes it easier for us to incorporate a wide range of color parameters in the training data, uniformly sample from the RGB space. However, since we expect many metallic objects to have commonly seen colors of golden or silver tones and want our network to represent these colors well, we took care to include a larger number of material samples with common metal colors in the training.

Given a reflection color c_{ref} , we use a constant edge color $c_{\text{edge}} = (1, 1, 1)$ and derive the complex-valued IOR of the metallic object. We acknowledge that using this conversion can lead to unrealistic IOR values if the input color is highly saturated—e.g. leading to $\eta_r = 1.0$ and $\eta_i = 0.0$ for one of the RGB channels. However, using true Fresnel coefficients computed from complex-valued IORs, instead of an empirical Fresnel calculation where a color is used as the specular albedo and blended toward white, gives us a much more realistic dip in reflection intensity at around 80° incident angle.

To summarize, parameters of the conductive core GGX lobe include:

- RGB-valued conductor IOR, real part (η_r), derived from reflection color (c_{refl})
- RGB-valued conductor IOR, imaginary part (η_i), derived from c_{refl}
- Lobe roughness α_{core} , identical to that of the dielectric core lobe

Conductive Haze Lobe (ID: 4)

Similar to the dielectric haze lobe, this is again a rougher companion of the conductor core lobe, where only the roughness α_{haze} is different. The conductive haze lobe is associated with the following parameters:

- Conductor IOR (η_r, η_i), identical to that of the conductive core lobe
- Lobe roughness (α_{haze}), identical to that of the dielectric haze lobe

In addition to the lobe parameters, the mixing factor (m_{ch}) between the core and haze lobes—which is the same for the dielectric and conductive branches—is another parameter we use in our model. This mixing factor modulates the shape of the reflection distribution. If $m_{\text{ch}} = 0$ or $m_{\text{ch}} = 1$, we revert to a single-GGX distribution. A second mixing factor m_{dc} adjusts the amount of dielectric or conductor reflection. We prefer values of m_{dc} close to 0 or 1, since real-world materials are either dielectrics or conductors and hardly ever exhibit a completely mixed state (semi-metallic).

Symbol	Range	Distribution
η_d	(1.0, 2.0)	$\mathcal{N}(1.5, 0.2)$
η_r	—*	—*
η_i	—*	—*
α_{core}	(0.001, 0.2)	$\mathcal{S}(0, \sqrt{0.2})$
α_{haze}	(0.005, 0.5)	$\alpha_{\text{core}} + \mathcal{U}(0, 0.3)$
m_{ch}	(0.0, 1.0)	$\mathcal{N}(0.5, 0.2)$
m_{dc}	(0.0, 1.0)	mainly around 0 or 1

Table 1: Free parameters describing the core and haze lobes and the distributions they follow when we construct our universal MLP training data. $\mathcal{S}(a, b)$ refers to the distribution achieved from taking the square of values drawn from $\mathcal{U}(a, b)$, favoring smaller values.

*The distributions of η_r and η_i are not straightforward to describe mathematically, since these parameters are derived from an RGB color, c_{refl} . In our training, this color is sampled following $\mathcal{U}(0.0, 1.0)^3$ and supplemented with a dedicated set of metallic-hue (e.g. golden, silver) samples.

For our per-material uplifting (as opposed to generating data for universal MLP training), we choose to apply a power function to the original roughness texture to achieve core and haze roughness textures, as mentioned in Eq. 1 from the main paper and Section 2 in this document. Fig. 2 demonstrates the double reflection highlight effects on the fruit from our paper’s teaser image.

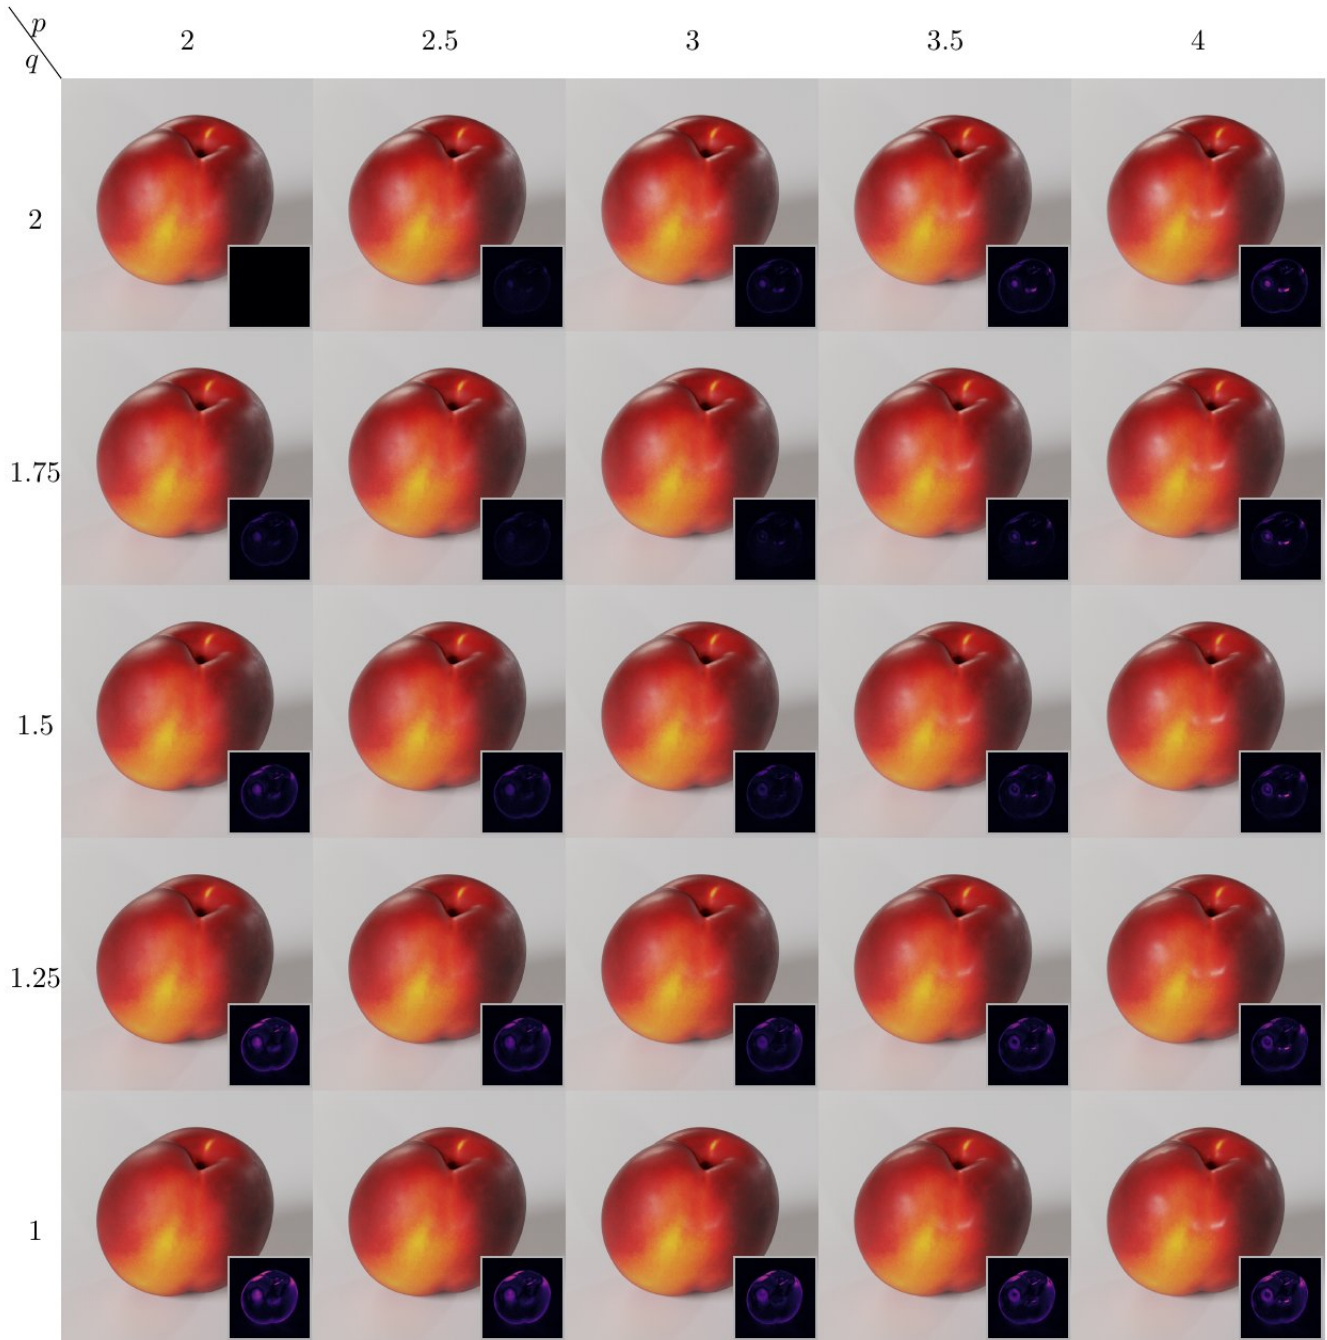


Figure 2: 5×5 grid of haze images. We increase p and decrease q to split a single roughness value into two values, allowing continued use of the original roughness texture. The mixing weight m_{ch_d} is set to 0.8 in this example. The average roughness α of the input texture is 0.388. Setting $p = q = 2$ reproduces the original roughness map. The TLIP image in the lower right corner shows the difference to the original input without the double highlight feature.



Figure 3: By increasing the thickness parameter t_{dust} , the object becomes more fuzzy. In this example, we increase t_{dust} from 0 to 0.6 in equal steps. The FLIP image in the lower right corner shows the difference to the original material without fuzz.

Dust/Fuzz Lobe (ID: 1)

This is a sheen lobe [Kulla and Conty, 2017] we use to simulate thin, retro-reflective and strongly view-dependent behavior, as can be observed from surfaces covered with dust or fine and short hairs (fuzz). When this lobe is enabled, the surface appears more opaque when viewed from grazing angles, as compared to near-normal directions. This lobe also weakens the intensity of the specular lobes.

The dust/fuzz lobe is controlled by the following parameters:

- Lobe weight (w_{dust}), the overall scale of the lobe, representing coverage. With $w_{\text{dust}} < 1.0$, the view-dependent appearance becomes less salient. We disable the lobe by setting $w_{\text{dust}} = 0$
- Dust thickness (t_{dust}), the thickness of the dust/fuzz layer. If the thickness is small, the lobe barely affects the appearance at near-normal angles while still creating an opaque appearance when viewed from grazing. If the thickness is high, normal and grazing angle appearances become more similar. Note that this parameter is traditionally called “roughness,” but we believe the term “thickness” illustrates the concept better and use it throughout our texts. Figure 3 demonstrates the effect of the parameter.

In our work, we fix the scatter color of this lobe c_{dust} to be white. To create more diverse enhanced appearances, this color parameter can be adjusted to simulate dirt or rust.

Clearcoat Lobe (ID: 2)

The clearcoat lobe is an almost perfectly specular dielectric GGX lobe that simulates effects like glazing or clear paint on top of another material. Unlike the dust/fuzz lobe, the weight w_{clear} of this lobe is binary, indicating the coating is either present or not. The IOR η_{clear} is an independent parameter, since the coating represents a different medium than that of the main dielectric lobes. The GGX roughness of the clearcoat lobe α_{clear} is fixed as 0.005.

This lobe is controlled by the following free parameters:

- Binary lobe weight (w_{clear}), either 0 or 1. $w_{\text{clear}} = 1$ indicates the presence of the clearcoat lobe, and $w_{\text{clear}} = 0$ disables it



Figure 4: Clearcoat adds a second, almost perfectly specular highlight on top of the object. As the IOR η_{clear} is increased, the object becomes progressively more reflective, altering the perceived nature of the coat. In this example we increase η_{clear} from 1.3 to 1.7. The FLIP image in the lower right corner shows the difference to the simple PBR input without clearcoat.

- Clearcoat IOR (η_{clear}), assuming a different range than the dielectric core/haze IOR and chosen independently (see Fig. 4 for an example)

Symbol	Range	Distribution
w_{dust}	(0.0, 1.0)	$\mathcal{U}(0.0, 1.0)$
t_{dust}	(0.0, 1.0)	$\mathcal{U}(0.0, 1.0)$
w_{clear}	{0.0, 1.0}	Binary
η_{clear}	[1.3, 1.7]	$\mathcal{N}(1.5, 0.08)$

Table 2: Free parameters describing the dust and clearcoat lobes (above the mandatory lobes), along with the distributions they follow when we construct our universal MLP training data.

Inner Scatter Lobe (ID: 7)

This lobe is another sheen lobe but placed *below* the specular lobes. This lobe simulates scattering within thin layers like latex, wax or gel that are angle-dependent and cannot be modeled with diffuse lobes. The visual appearance is significantly different than that of the dust lobe, since the Fresnel effect of the specular lobes will reduce the strong grazing angle effects the sheen lobe would otherwise have. Overall, this lobe results in an appearance that resembles strong scattering of thin media since the retro-reflectivity is retained. An example of the lobe can be seen in Fig 5.

The inner scatter lobe is described with the following parameters:

- Lobe weight (w_{inner}), the overall scale of the lobe. If $w_{\text{inner}} = 0$, the lobe is disabled
- Inner scatter thickness (t_{inner}), similar to the thickness of the dust lobe but resulting in more uniform appearance (see above)
- Scatter color (c_{sca}), color of the lobe (i.e. the light that gets not absorbed but scattered), which is multiplied into the reflectance



Figure 5: The inner-scatter lobe adds a scatter-like component beneath the surface highlight. Unlike the fuzz or dust lobe, it can take on arbitrary colors, since the scatter color is not necessarily tied to the diffuse color. In this example, we progressively increase the weight w_{inner} from 0 to 1. We use the same red-colored hue as in the paper’s teaser image. The ∇ LIP image in the lower right corner shows the difference to the original input without inner scatter.



Figure 6: The subcutaneous scatter lobe adds a third specular highlight beneath the specular highlights from the surface. Depending on the roughness of α_{haze} , this highlight can become very subtle. Unlike the dielectric surface highlights, which remain uncolored, the subcutaneous scatter lobe can be tinted. The two images illustrate the difference with the lobe enabled and disabled. The ∇ LIP image in the lower right corner shows the difference between subcutaneous scatter on and off.

Symbol	Range	Distribution
w_{inner}	(0.0, 1.0)	$\mathcal{U}(0.0, 1.0)$
t_{inner}	(0.0, 1.0)	$\mathcal{U}(0.0, 1.0)$
c_{sca}	$[0.0, 1.0]^3$	$\mathcal{U}(0.0, 1.0)^3$
w_{subc}	{0.0, 1.0}	Binary
α_{subc}	(0.1, 0.95)	$\alpha_{\text{haze}} + \mathcal{U}(0.05, 0.45)$

Table 3: Free parameters describing the inner and subcutaneous scattering lobes (below main lobes), along with and the distributions they follow when we construct our universal MLP training data.

Symbol	# parameters	mapping
η_{d_core}	1	η_d
η_{d_haze}	1	η_d
m_{ch_d}	1	m_{ch}
α_{d_core}	1	α_{core}
α_{d_haze}	1	α_{haze}
η_{r_core}	3	η_r
η_{i_core}	3	η_i
η_{r_haze}	3	η_r
η_{i_haze}	3	η_i
m_{ch_c}	1	m_{ch}
α_{c_core}	1	α_{core}
α_{c_haze}	1	α_{haze}
m_{dc}	1	m_{dc}
w_{dust}	1	w_{dust}
t_{dust}	1	t_{dust}
c_{dust}	3	(1,1,1)
w_{clear}	1	w_{clear}
η_{clear}	1	η_{clear}
α_{clear}	1	0.005
w_{inner}	1	w_{inner}
t_{inner}	1	t_{inner}
c_{inner}	3	c_{sca}
w_{subc}	1	w_{subc}
η_{subc}	1	η_d
α_{subc}	1	α_{subc}
c_{subc}	3	c_{sca}

Table 4: Comparison of the full set of input parameters of the model vs. the simplified coupled version.

Subcutaneous Scatter Lobe (ID: 8)

The subcutaneous scatter lobe is, like the inner scatter lobe, placed below the specular lobes. Unlike the inner scatter lobe, it is a rough GGX lobe that has a tint color multiplied to its reflectance. It simulates light that enters a medium and goes through a small number of scattering events before exiting again. The visual result is a soft, broader and colored highlight added to the sometimes hard and synthetic appearing uncolored specular lobe, which can give a material a softer feeling. A real-life example that exhibits such a behavior is leather.

The lobe shares its IOR with the specular lobes, i.e. its IOR is given by η_d . Moreover, while the scatter color c_{subc} is in theory a parameter that can be chosen independently from the inner scatter lobe, in our experiments we adopt the same scatter color c_{sca} for both lobes. We use this lobe in an empirical fashion to modulate the appearance of a thin scattering layer. Fig. 6 shows the effect of the lobe.

The subcutaneous scattering lobe is determined by the following free parameters (along with preset ones

mentioned above):

- Binary lobe weight (w_{subc}), for enabling and disabling the lobe
- Lobe roughness (α_{subc}), since light is consecutively roughened when it passes through different media, the roughness is higher than the roughness of the specular lobes
- Scatter color (c_{sca}) determines the tint of the highlight. The color is multiplied *after* the evaluation of the Fresnel coefficient.

Note that there are some inter-dependencies in certain lobe parameters, as we will further describe in Section 2. As already mentioned, we couple parameters like IORs and roughness values of different lobes. With our lobe design and parameter assignment scheme, we need 22 free parameters to fully define our multi-lobe BRDF (if we do not count the parameters we hard-code to a fixed value). All our lobe parameters and relevant free parameters are shown in Table 4.

2 Enhancement Types

Having introduced each lobe in our enhanced non-diffuse BRDF model, we now describe our six enhancement types from the main paper in more detail. We recall from the main paper that our baseline (Type 0) enhancement consists of the four main specular reflection lobes discussed in Section 1, and other enhancements (Type 1–5) involves one or more decorator lobes.

Type 0: Haze (ID: 3, 4, 5, 6)

This baseline enhancement from a simple PBR material can be summarized as constructing our four specular reflection lobes (with spatially varying parameters) from the simple PBR material’s base color, roughness, and metalness textures, as described in the main paper.

At each texel, we need to determine the parameters η_d , m_{ch} , α_{core} , α_{haze} , η_r , η_i , and m_{dc} .

- The dielectric IOR (η_d): This parameter is chosen to be constant per-material, assuming the material has a fixed type of underlying medium. η_d is sampled from $\mathcal{N}(1.5, 0.2)$, matching the distribution reported in Section 1. If an asset contains multiple materials, each material’s dielectric IOR is sampled independently from this Gaussian distribution.
- Core/haze mixing factor (m_{ch}): While this parameter can land anywhere in $[0, 1]$, in our enhancement experiments, we chose to sample it from $\mathcal{U}(0.45, 0.55)$ for each material, so that both lobes have significant contributions. In other words, a mixing factor close to 0 or 1 emphasizes only one of the two lobes, making the material fall back to a simple PBR material (if no decorator lobe is present). Here, we are more interested in creating appearances with double specular highlights.
- The core roughness (α_{core}): As described in the paper, we use a simple model that computes this roughness from the input PBR roughness; the specific function is chosen as $\alpha_{\text{core}} = r^p$, where p is sampled from $\mathcal{U}(2.8, 3.2)$. With this mapping function, unless the input roughness r is very

high, α_{core} lands in its supported range in the training data. Even when α_{core} lands outside of $(0.001, 0.2)$, in most of our experiments the neural reconstruction of the material does not seem to be very significantly affected.

- The haze roughness (α_{haze}): We use another function to compute this parameter from the input PBR roughness, $\alpha_{\text{haze}} = s \cdot r^q$, where s is sampled from $\mathcal{U}(1.2, 1.6)$ and q is sampled from $\mathcal{U}(1.4, 1.6)$. The same argument about the global range of this parameter seems to apply here as for α_{core} .
- The complex-valued conductor IOR (η_r, η_i): As described in Section 1, these values are derived from a specular reflection color, which we choose to be our PBR base color.
- Dielectric/conductive mixing factor (m_{dc}): Intuitively, we use the PBR metalness parameter unmodified as m_{dc} . Note that metalness values close to 0.5 is represented less than values closer to 0 or 1 in our training data, but are also not entirely unrepresented.



Figure 7: Examples of Type 0 material augmentations, compared to their PBR equivalents.

To summarize, we leverage our PBR parameters and some randomness to derive a set of parameters for the four specular lobes in our model, intentionally creating nearly evenly mixed core and haze lobes with contrasting roughness scales that can lead to double-highlight appearances on materials. This core–haze construction is applied to all material instances. If no additional decorator lobes are enabled, this concludes our Type 0 enhancement. Fig. 7 shows some examples of Type 0 enhanced materials.

Type 1: Dust (ID: 1)

In this enhancement type, we enable the dust/fuzz lobe by setting its weight nonzero. In particular, we need to additionally sample or compute the (spatially varying) parameters w_{dust} and t_{dust} .

- Dust weight (w_{dust}): To realistically simulate dust accumulation, we drive the weight of the dust lobe in a material using the world-space normal vectors of the underlying asset—naturally, with higher weight on upward-facing surfaces. We also infuse spatial variations in dust coverage, such as fingerprints or wiped patterns, using pre-selected masks shown in the main paper. We randomly select from a pool of seven different textures with the features we want to simulate (low and high coverage, wipes, fingerprints, two types of “wet” dust with water ripples, and a uniform distribution). These textures are mixed with a constant value as “base” coverage of the dust and further multiplied with the weight determined from the geometric normal. The base coverage value is a scalar in the range $(0, 1)$, and we select smaller values more frequently since we want to preserve the visual breakup. For 20% of the data we chose a minimum value of 0.5 to guarantee that enough exemplars have high coverage. Since all our textures are periodic and seamless, we can apply them with tri-planar mapping. The dust weight textures are computed in Blender with custom Python scripts.
- Dust thickness (t_{dust}): This parameter is chosen to be constant per-material, sampled from $\mathcal{U}(0, 1)$ independently for each material, reflecting the distribution in the universal training data.




Figure 8: Examples of Type 1 material augmentations, compared to their PBR equivalents.

Fig. 8 shows some examples of Type 1 enhanced materials.

Type 2: Clearcoat (ID: 2)

We enable the clearcoat lobe by setting its weight to 1 (almost everywhere). Formally, we need to determine the (spatially varying) parameters w_{clear} and η_{clear} .

- Clearcoat weight (w_{clear}): Each material associated with an asset can be a constant material, or could be a textured one with spatial variations. With a constant material, we assign its clearcoat weight to 1 with a probability of 70%. This enables us to create assets which are partially or fully coated, improving the diversity in our enhanced materials. For textured materials, we associate its clearcoat weight with a binary mask that simulates damages such as scratches or dents on the coating. Our pre-selected binary weight masks are periodic, and for the same asset and texture space, using a cropped subregion of a mask introduces large-scale damages on the coating, while tiling multiple copies of the mask and using it as a whole leads to small-scale imperfections. We heuristically determine the scale of the binary masks based on the bounding box size of each object, but since the object meshes may not reflect real-world sizes, the quality of results vary with assets (still, our experimented heuristic infused most materials with appropriately sized visual breakups). We use three different clearcoat damage textures (scratches and two forms of dents).
- Clearcoat IOR (η_{clear}): The IOR parameter of each material is independently sampled from $\mathcal{N}(1.5, 0.08)$, matching the distribution in the training data.

Fig. 9 shows some examples of Type 2 enhanced materials.




Figure 9: Examples of Type 2 material augmentations, compared to their PBR equivalents.

Type 3: Dust + Clearcoat (ID: 1, 2)

This is a simple combination of Type 1 and Type 2, producing surfaces that exhibit both a glossy coating and accumulated dust. Parameter choices of each lobe are made as for the two previous types. Fig. 10 shows some examples of Type 3 enhanced materials.




Figure 10: Examples of Type 3 material augmentations, compared to their PBR equivalents.

Type 4: Fuzz + Inner Scatter + Subcutaneous Scatter (ID: 1, 7, 8)

For this enhancement type, we apply the dust/fuzz lobe with a uniform weight over dielectric objects to produce a fuzzy appearance, and also enable the inner scatter and subcutaneous scatter lobes in the dielectric branch. The scatter color, shared by both scattering lobes, is derived from the base color via gentle desaturation and hue shifting. More specifically, the parameters w_{dust} , t_{dust} , w_{inner} , t_{inner} , c_{sca} , w_{subc} , and η_{subc} are chosen as follows:

- Weight of the dust/fuzz lobe (w_{dust}): When the top-layer sheen lobe is used to create a fuzzy, rather than dusty appearance, the weight parameter is chosen as constant per-material. To make the fuzzy appearance salient, we sample the lobe weight from $\mathcal{U}(0.5, 1.0)$ for each material (avoiding overly small weight).
- Thickness of the dust/fuzz lobe (t_{dust}): This parameter is sampled from $\mathcal{U}(0.0, 1.0)$ for each material.
- Weight of the inner scatter lobe (w_{inner}): Similar to the weight of the fuzz lobe, we sample this parameter from $\mathcal{U}(0.5, 1.0)$ for each material.

- Thickness of the inner scatter lobe (t_{inner}): As a pure design choice, we sample this parameter from $\mathcal{U}(0.3, 0.4)$ for each material. This choice encourages a desired amount of contrast between appearances at near-normal and grazing viewing angles.
- The scatter color (c_{sca}): We set this color parameter to be loosely coupled to our input PBR base color. For stylistic purposes, we first perform a hue-shift on the original base color (keeping saturation and value unchanged and offset the hue by an amount of 0.1), and then slightly desaturate the color by scaling the RGB color from $[0.0, 1.0]^3$ to the range of $[0.2, 1.0]^3$. Note that this assignment of the scatter color is one of the many design choices one could make, and since our universal neural material training data has incorporated a wide range of scatter colors, most materials with our chosen scatter colors can be reliably reconstructed.
- The weight of the subcutaneous scatter lobe (w_{subc}): As this weight is binary, enabling the lobe requires setting $w_{\text{subc}} = 1$ in each material.
- The subcutaneous scatter roughness (η_{subc}): This parameter is computed as the sum of η_{haze} and an offset sampled from $\mathcal{U}(0.05, 0.45)$, as how it is chosen per-material in the training data.




Figure 11: Examples of Type 4 material augmentations, compared to their PBR equivalents.

Fig. 11 shows some examples of Type 4 enhanced materials.

Type 5: Fuzz + Clearcoat + Inner Scatter + Subcutaneous Scatter (ID: 1, 2, 7, 8)

This last enhancement type enables all the decorator lobes. If the fuzz weight is high, the fuzz reduces the visual sharpness of the clearcoat, often resulting in a milky or gel-covered appearance. Parameters related

to the fuzz, inner scatter and subcutaneous scatter lobes are chosen as in Type 4, while the clearcoat parameters are chosen as for Type 2. However, to model smooth, milky appearance on assets, we did not introduce damages on the clearcoat as in Type 2—all coatings are modeled as smooth and intact. Figure 12 shows some examples of Type 5 enhanced materials.




Figure 12: Examples of Type 5 material augmentations, compared to their PBR equivalents.

Please note that as mentioned with multiple enhancement types, some parameter choices reflect our design decisions and our uplifted appearances do not span the full range of appearances supported by our model or covered in our training data. This is because our goal in this current project is to introduce a few types of interesting appearances that cannot be achieved with simple PBR models, rather than exploring the full range of materials specified with a new model.

3 Neural Training

We now expand on Section 4 in our main paper, providing more details on our neural architecture and training specifics, training data for our universal MLP, and our per-material neural reconstruction quality.

3.1 Network Architecture

Following the architecture of Zeltner et al. [2024], we train two separate universal MLPs for BSDF evaluation and sampling. Both MLPs are designed to decode the same latent textures, with the evaluation MLP producing BSDF values and sampling MLP producing parameters used for BSDF importance sampling. Similarly, we additionally use a third MLP to predict the BSDF’s transmission albedo and luminance

Stage	Size	Iterations	Loss	Timing
Universal BRDF	4×64	500k	ReLU2	1051.83 s
Universal Sampler	2×32	100k	KL	426.83 s
Universal Albedo	4×64	500k	L1	843.42 s
Per-Material Optimization	–	100k	L1	380.64 s

Table 5: Summary of the network architectures, training schedules, loss functions, and runtimes for each stage of our pipeline. Per-material optimization runtime is averaged over a fraction of our examples.

reflectance. Due to limited space in the main paper, we did not explicitly mention that BSDF values and albedo information are in fact produced by different MLPs—they can also be produced by the same MLP with more output channels, if desired.

We kept the universal MLPs shallow to reduce their evaluation time and used a low-dimensional latent space to limit the number of channels/parameters involved in our downstream generative process. The final architecture (see Table 5) and dimensionality (6D) were chosen empirically. Adding additional hidden layers yielded diminishing returns, and four latent dimensions was the minimum number of dimensions required to represent color information, while six dimensions represent the optimal balance that allowed for reliable reconstruction over a wide range of materials.

For both our universal MLP training and later-stage, per-material latent optimization, we use an Adam optimizer with a learning rate of 0.001 ($\beta_1 = 0.9$, $\beta_2 = 0.999$). We use stratified sampling on the reference material data and a sampling scheme that combines Rusinkiewicz sampling with importance sampling of the reference BSDF in a ratio of 90% : 10%. Our universal training was performed on a single RTX 5090 GPU, and per-material optimizations were performed on a multi-user GPU cluster with L40 and L40S GPUs, as mentioned in the paper. Table 5 summarizes the network sizes, losses, iteration counts, and training time.

3.2 Universal Training Data

We now describe how we generate data for training our universal MLPs mentioned in the previous section. Our goal is to use the data to guide the universal MLPs’ limited capacity toward the most important features and aspects of materials represented with our multi-lobe model—promoting learning where it matters the most.

Parameter Layout

Our dataset is generated procedurally from random inputs, which are mapped to material parameters by sampling from our predefined parameter distributions. This biases generation toward materials that occur with higher frequency, ensuring that common materials are represented more frequently. For example, we sample η_d from $\mathcal{N}(1.5, 0.2)$, with the highest density at 1.5, since this is the estimated IOR distribution in real-world dielectric materials.

Because our data generation is procedural, our training data is essentially unlimited and does not come with any memory overhead. For training, we arrange input data into five 8192×8192 data grids. We

partition four of these grids into 128×128 patches, which we treat as individual “training materials.” Within these patches, only the roughness of the core specular lobes α_{core} and the mixing factor between the core and haze lobes m_{ch} vary spatially. We increase α_{core} linearly along the x -axis and m_{ch} along the y -axis of the grid, and choose a non-uniform distribution centered around 0.5 for m_{ch} . The haze roughness α_{haze} is randomized but coupled to α_{core} by adding a per-patch constant offset

$$\alpha_{\text{haze}} = \alpha_{\text{core}} + \Delta_{\text{haze}} \quad (\text{S.7})$$

where Δ_{haze} is randomly sampled. This roughness variation ensures sufficient intra-material diversity. Since many simple PBR materials—which we use as inputs to our material enhancement—include roughness textures, we specifically target a parameter distribution that captures small-scale input variations.

Additionally, within each 8192×8192 data grids, we reserved three 2048×2048 regions (still partitioned into 128×128 patches) for representing color variation. We used multiple HSV color ramps with varying hue and value, and seven discrete saturation levels. Each chosen reflection color c_{refl} is then used to derive η_r and η_i for the conductor lobes, as well as c_{sca} for the inner and subcutaneous scatter lobes. All parameters other than roughness, core-haze mixing factor, and reflection/scattering colors are then chosen randomly but held constant within each 128×128 patch. These patched data grids are what we call the four *themed* data classes, as in the main paper.

Finally, a fifth of the training set is generated without bias toward material features that are important to represent. For each material, all parameters are initialized randomly. Specifically, every parameter is sampled independently and uniformly within its allowed range. We refer to the resulting data as the *infill* class, and use it to encourage a more uniform latent space population, which benefits the downstream generative process. A visualization of all our training data is in Fig. 13.

We also visualize the completely random training data mentioned in Section 4.2 of the paper. In this setting, as in the *infill* class, every cell of each 8192×8192 grid uses an independent set of parameters, chosen from uniform distributions within respective ranges.

More on Themed Classes

As mentioned in the paper and shown in Fig. 13, we have four themed classes which we call *clearcoat*, *dust*, *haze* and *scatter*. The *haze* class has no decorator lobe and uses procedural initialization on all patches. In each of the other themed classes, the corresponding decorator lobe is always enabled—lobe 1 for *dust*, lobe 2 for *clearcoat* and lobe 7 and 8 for *scatter*), while the remaining lobes are assigned randomized parameters, including weights that enable or disable these lobes. This randomization follows additional priors, such as 20% of the patches include an active clearcoat lobe and 10% include an active dust lobe. In some patches, we hold subsets of parameter combinations fixed and vary only a single parameter, to help the network disentangle the effects of individual parameters. We introduce these themes to explicitly expose the network to the types of material enhancements we aim to model. The conductor lobes are enabled in all themes except *scatter*, because the scatter lobes are only relevant to the dielectric branch, as volumetric scattering typically does not occur in conductors.

Lastly, outside of the regions dedicated to color variation (mentioned previously), the specular reflection color c_{refl} used to calculate η_r and η_i of the conductors is sampled from two color pools: a set of fully random colors distributed uniformly in $[0, 1]^3$ and a set of common metal colors. We sample from the




Figure 13: Top: A visualization of our training data. In our informed training, we have four themed classes and a random-parameter infill class; counterclockwise: scatter, clearcoat, infill, dust, haze. Bottom: Completely randomizing the parameters leads to a worse performance of the universal MLP.

random-color pool with a probability of 20% and sample uniformly from the common-color pool otherwise. Finally, 30% of the patches are conductor-dominant and 70% are dielectric-dominant.

3.3 Per-Material Fitting Quality

To assess the ability of the universal network to fit our enhanced materials and evaluate the per-material optimization quality, we randomly selected 50 examples from each enhancement type and reconstructed each analytical material using the neural representation, by optimizing the per-material neural latents against the fixed universal MLPs, using the combined loss function introduced in Section 4.3 of our paper. We then rendered the reference and reconstructed materials on preview spheres, computed \mathcal{F} LIP [Andersson et al., 2020] error images, and summarized each reconstruction by its mean \mathcal{F} LIP error. Table 6 reports the best, median, and worst mean \mathcal{F} LIP values across the evaluated materials per type.

Type	Best	Medium	Worst
Type0	0.0487	0.0920	0.1871
Type1	0.0312	0.0727	0.1670
Type2	0.0649	0.0858	0.1192
Type3	0.0554	0.0877	0.1664
Type4	0.0276	0.0643	0.1705
Type5	0.0454	0.0807	0.1754


Table 6: Best, medium, and worst \mathcal{F} LIP error for each type.

Moreover, we visualize the reconstruction quality in a few ways in Fig. 14–19, showing example \mathcal{F} LIP error images from each enhancement type, some worst-quality reconstruction results, and the mean \mathcal{F} LIP error histograms. These illustrations show that inaccuracy in our neural reconstruction typically manifests as color shifts and modified highlight shapes. Most fitting errors are acceptable, as our enhanced materials, if passing our furnace test, would simply exhibit a slightly different scatter color or core/haze roughness. Still, the errors can occasionally become quite pronounced, as exemplified in Fig. 18. We believe such pronounced error in this particular example (and similar ones) was due to the large error in transmission albedo dominating our combined fitting loss, compromising the accuracy in BRDF.

4 Diffusion-Generated Materials

We show rendered images of generated neural materials for our full test set of 3D models in Fig. 20. While results are generally encouraging, we note that generative modeling of neural textures has remaining challenges. First, as detailed in Section 4.2 of the main paper, the diffusion model is trained on latent space MSE loss without knowledge of the constraints (e.g. valid latent space regions) of the neural material. In other words, there is no easy way to strictly enforce generating valid latents. Second, the generated views of the neural latents are projected to texture space, which assumes that the diffusion model respects the input geometry, and that the neural textures can be filtered and interpolated. We observe sporadic fireflies (e.g. the Volkswagen Beetle car model) in regions with missing view coverage, and this issue is greatly amplified when rendering with global illumination. In this section, we provide

(a) Example FLIP images: Type 0



(b) Worst-case reconstruction



(c) Mean FLIP histogram





Figure 14: Visualization of the neural reconstruction quality in Type 0 materials. (a) Example FLIP images from six reconstructed materials, with varying neural fitting qualities. In some images, the FLIP error intensity appears significantly different on the left and right sides of the preview ball—this is because we tiled up multiple material textures and performed neural fitting in the same pass, and these sub-materials have different fitting qualities. (b) Two materials with the worst reconstruction qualities from our test set of 50. The first example exhibits a purple hue shift, and the second example demonstrates both a slight color drift and modified highlight shapes. (c) A histogram of the mean FLIP values in our 50 error images.

(a) Example FLIP images: Type 1



(b) Worst-case reconstruction



(c) Mean FLIP histogram

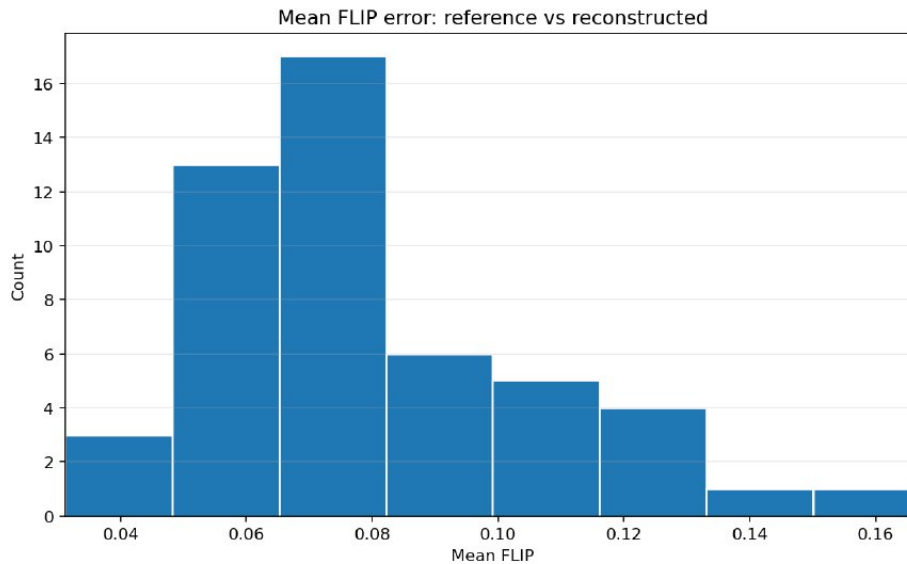



Figure 15: Visualization of the neural reconstruction quality in Type 1 materials. (a) Example FLIP images from six reconstructed materials, with varying neural fitting qualities. (b) Two materials with the worst reconstruction qualities from our test set of 50. Both examples exhibit some color drift. Note that the preview ball was merely a default object used to demonstrate our textured materials, and the dust layer may not be at the top—for instance, the dusted area forms a band on the preview ball. (c) A histogram of the mean FLIP values in our 50 error images.

(a) Example FLIP images: Type 2



(b) Worst-case reconstruction



(c) Mean FLIP histogram





Figure 16: Visualization of the neural reconstruction quality in Type 2 materials. (a) Example FLIP images from six reconstructed materials, with varying neural fitting qualities. (b) Two materials with the worst reconstruction qualities from our test set of 50. Overall FLIP error seems to be lower than other types, and color and highlight differences in the reference and fitted materials are not extremely disturbing. (c) A histogram of the mean FLIP values in our 50 error images.

(a) Example FLIP images: Type 3



(b) Worst-case reconstruction



(c) Mean FLIP histogram





Figure 17: Visualization of the neural reconstruction quality in Type 3 materials. (a) Example FLIP images from six reconstructed materials, with varying neural fitting qualities. (b) Two materials with the worst reconstruction qualities from our test set of 50. (c) A histogram of the mean FLIP values in our 50 error images.

(a) Example FLIP images: Type 4



(b) Worst-case reconstruction



(c) Mean FLIP histogram






Figure 18: Visualization of the neural reconstruction quality in Type 4 materials. (a) Example FLIP images from six reconstructed materials, with varying neural fitting qualities. (b) Two materials with the worst reconstruction qualities from our test set of 50. The error in the scatter color in the second example is very pronounced; we found that this is due to the error in the transmission albedo dominated our combined fitting loss (despite only assigned a weight of 0.04), compromising the BRDF reconstruction quality. (c) A histogram of the mean FLIP values in our 50 error images.

(a) Example FLIP images: Type 5



(b) Worst-case reconstruction



(c) Mean FLIP histogram

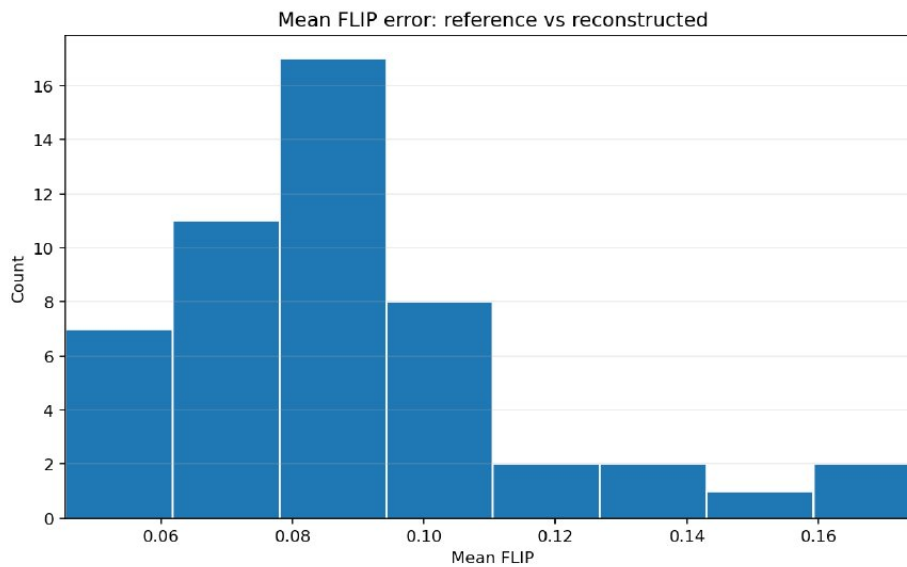


Figure 19: Visualization of the neural reconstruction quality in Type 5 materials. (a) Example FLIP images from six reconstructed materials, with varying neural fitting qualities. (b) Two materials with the worst reconstruction qualities from our test set of 50. Both examples exhibit a similar color drift (pastel yellow becoming pastel blue) resulting from the reconstruction. (c) A histogram of the mean FLIP values in our 50 error images.

more detail into how we evaluate our generative model, as well as additional information that we did not get to cover in the main paper.

4.1 Evaluation details

For both our method and the compared one (TRELLIS.2), we assume known 3D geometry in each asset and focus on material generation from text or a single image. For our test set (Fig. 20) with 16 3D objects, each with a reference neural material, we render reference images for 16 views under controlled lighting. These views are used to create descriptive text captions using the visual language model (VLM) Qwen2.5-VL-7B. We also add explicit keywords specific to enhancement types to the captions.

To evaluate the pipeline, we present three experiments below. Our baseline is TRELLIS.2 [Xiang et al. \[2025\]](#), a recent image-guided PBR material generation method. We include two variants of our generation pipeline, differing only in which base color texture is used.

TRELLIS.2 The TRELLIS.2 texture generation pipeline generates PBR materials from a 3D mesh and an example image input. For each example, we use the mesh from the test set and a reference image (as described above) as the example image. The method produces a 3D mesh with PBR material (base color, roughness, and metalness).

Our_{trellis} We use the base color texture generated by TRELLIS.2, the mesh and the caption (generated from reference images, as described above) as input and produce a 3D model with a material consisting of the TRELLIS.2 base color and a generated neural specular component.

Our_{ref} We use the base color texture available with our reference material, the mesh and caption as input, and produce a 3D model with a material consisting of the reference base color and a neural non-diffuse component.

We generate five random seeds of each asset and subjectively pick the best result. For each experiment we render 16 views of the object shaded with the generated material, with identical camera poses and lighting condition to the reference images. We then compute error metrics for quantitative evaluation.

The first experiment establishes a baseline by testing how well TRELLIS.2 can generate materials from a single example image. The Our_{trellis} experiment determines how well our model can enhance the TRELLIS.2 material by adding a neural specular BSDF lobe. Finally, the Our_{ref} experiment shows how well our method can reconstruct the specular BSDF given a high quality artist created base color and descriptive prompt. We are aware of that these pipelines perform slightly different tasks, but we hope to provide a reasonable evaluation.

Fig. 21–24 show our test objects rendered with materials in four variants: the reference neural materials, PBR materials generated from TRELLIS.2, base color from TRELLIS.2 combined with neural, non-diffuse materials generated from our method, and known base color combined with our generated neural materials. Comparison between the Our_{trellis} and Our_{ref} variants show that while base colors are left unmodified, the specular enhancements produce interesting and richer visual effects in most scenes.




Figure 20: Generated neural materials for all 3D models in our test set. Intended enhancement types of these assets' reference neural materials are haze (a), dust (b, c), clearcoat (d, e, f, g), dust + clearcoat (h, i), fuzz + scatter (j, k, l, m), all lobes (n, o, p).



Figure 21: Our test set objects (a)–(d) rendered with materials in four variants: the reference, hand-crafted neural materials, PBR materials generated from TRELLIS.2, base color from TRELLIS.2 combined with neural, non-diffuse materials generated from our method, and known base color combined with our generated neural materials. Note that the Our_{ref} variant shows the same materials as presented in Fig. 20, while here we did not add ground planes and used black background. The artifacts in the car example was addressed as in the main paper.



Figure 22: Our test set objects (e)–(h) rendered with materials in four variants: the reference, hand-crafted neural materials, PBR materials generated from TRELIS.2, base color from TRELIS.2 combined with neural, non-diffuse materials generated from our method, and known base color combined with our generated neural materials. Note that the Our_{ref} variant shows the same materials as presented in Fig. 20, while here we did not add ground planes and used black background.



Figure 23: Our test set objects (i)–(l) rendered with materials in four variants: the reference, hand-crafted neural materials, PBR materials generated from TRELIS.2, base color from TRELIS.2 combined with neural, non-diffuse materials generated from our method, and known base color combined with our generated neural materials. Note that the Our_{ref} variant shows the same materials as presented in Fig. 20, while here we did not add ground planes and used black background.




Figure 24: Our test set objects (m)–(p) rendered with materials in four variants: the reference, hand-crafted neural materials, PBR materials generated from TRELIS.2, base color from TRELIS.2 combined with neural, non-diffuse materials generated from our method, and known base color combined with our generated neural materials. Note that the Our_{ref} variant shows the same materials as presented in Fig. 20, while here we did not add ground planes and used black background.

4.2 Raw prompts

Raw prompts for our test set, with the same ordering as appears in Fig. 20.

- (a) A silver Volkswagen Beetle rotates slowly against a black background, showcasing its iconic rounded body and classic design. The car's shiny surface reflects light, creating a hazy effect due to the reflective material. The camera captures the vehicle from various angles, highlighting its unique shape and vintage charm. A close-up shot from multiple perspectives. Haze.
- (b) The video showcases a luxurious urn with a uniform dust coating, rotating slowly against a black background. The urn is crafted from a speckled white marble, adorned with intricate gold detailing. The rim and base feature ornate gold bands, adding elegance to the piece. The camera captures the urn from various angles, highlighting its symmetrical design and the subtle reflections on the polished surfaces. The uniform dust coating enhances the texture and depth of the marble, creating a sophisticated appearance. A medium shot with a slight zoom-in effect. Uniform dust.
- (c) The video showcases a pair of dusty brown leather shoes with a textured surface, featuring silver buckles across the top. The shoes have a classic design with a rounded toe and a low heel. The camera rotates around the shoes, highlighting their worn condition with visible dust and scratches, emphasizing the texture and details of the leather. The background remains black throughout, keeping the focus entirely on the shoes. A close-up shot captures the heavy dust with ripples.
- (d) A rotating can of Pepsi, showcasing its sleek design and vibrant branding. The can is predominantly blue with white speckles, giving it a refreshing appearance. The iconic Pepsi logo is prominently displayed in orange and white, accompanied by the word "Pepsi" in large, bold letters. The can rotates smoothly against a black background, highlighting its glossy finish and the crisp, clean look of the label. The material appears to have an intact clear coat, ensuring the can maintains its pristine condition throughout the rotation. Intact clear coat.
- (e) The video showcases a beautifully crafted vase with a clear, intact coating. The vase has a flared top and a rounded base, both adorned with a subtle golden rim. The body of the vase is a smooth, translucent beige with a marbled texture, giving it a sophisticated and elegant appearance. The camera rotates slowly around the vase, highlighting its symmetrical design and the intricate details of its surface. The lighting accentuates the vase's glossy finish, creating a mesmerizing effect. A close-up shot captures the vase's intricate craftsmanship and the interplay of light and shadow across its surface. Intact clear coat.
- (f) A small, intricately designed teapot rotates against a black background. The teapot features a glossy, clear coating, giving it a translucent appearance. It is adorned with delicate blue floral patterns, adding a touch of elegance. The handle and spout are seamlessly integrated into the body, maintaining a cohesive design. The lid, slightly domed, complements the overall aesthetic. The teapot's surface reflects light, highlighting its smooth texture and intricate details. A close-up shot captures the teapot from various angles as it spins slowly. Intact clear coat.
- (g) The video showcases a round stool with a worn, metallic seat and a wooden frame. The seat appears to have a textured surface with visible scratches and marks, indicating damage to the clear coat.

The stool rotates slowly against a plain black background, highlighting its intricate design and the extent of the wear on the seat. The wooden legs are sturdy and ornately carved, adding to the stool's vintage aesthetic. A close-up shot captures the stool's details, emphasizing the contrast between the damaged seat and the intact wooden structure. Damaged clear coat.

- (h) A pomegranate with a uniform, smooth surface and intact clear coat rotates against a black background. The fruit's glossy, deep red skin reflects light, highlighting its natural texture and subtle variations in color. The stem is small and slightly curved, adding a touch of realism to the otherwise pristine appearance. The camera captures the pomegranate from multiple angles as it spins, emphasizing its spherical shape and the uniform surface. A close-up shot from various perspectives showcases the fruit's detailed texture and vibrant color. Uniform dust and intact clear coat.
- (i) The video showcases a beautifully crafted antique teapot with a shiny, metallic finish, likely silver or copper, exhibiting a rich patina. The teapot features a curved spout, two handles, and a lid with a small finial on top. It rotates smoothly against a black background, highlighting its intricate design and craftsmanship. The material appears uniform and intact, with a clear coat that enhances its polished surface. The camera captures the teapot from multiple angles, emphasizing its elegant form and historical charm. A close-up shot from various perspectives provides a detailed view of its ornate details and the subtle variations in its texture. Uniform dust and intact clear coat.
- (j) A table lamp with a beige fabric shade featuring subtle checkered patterns and a white base stands against a black background. The lamp is stationary, with no discernible movement or action taking place. The lighting highlights the texture of the fabric and the smoothness of the base. A static shot captures the lamp from a front-facing angle. Fuzz, with inner and subcutaneous scattering.
- (k) A 3D animated monkey character with a brown body and a lighter brown face is seen in a series of frames. The monkey is sitting on its haunches, holding its arms up near its face, and appears to be in motion, possibly dancing or moving rhythmically. The background is completely black, emphasizing the character's movements. The animation is smooth, with the monkey's limbs and tail swaying gently. The character's facial expression is neutral, and it has small, round ears and a short tail. The lighting is soft, highlighting the character's form against the dark backdrop. A medium shot capturing the character's full body. Fuzz, with inner and subcutaneous scattering.
- (l) A 3D model of a gray cat with a sleek, smooth texture and a long tail is shown against a black background. The cat walks gracefully, turning its head to look around, displaying a curious and alert expression. The animation captures the cat's fluid movements, with its paws moving in a natural walking pattern. The lighting highlights the cat's form, emphasizing its contours and the subtle details of its fur. A medium shot with a rotating camera perspective. Fuzz, with inner and subcutaneous scattering.
- (m) The video showcases a modern, minimalist chair with a tufted seat and backrest, upholstered in a soft, beige fabric that gives off a fuzzy, almost plush appearance. The chair features a sleek, metallic frame with a crossbar design for support. The camera rotates around the chair, highlighting its clean lines and ergonomic design. The background remains a solid black, emphasizing the chair's texture and form. The video captures the chair from multiple angles, providing a comprehensive view of its structure and comfort. Fuzz, with inner and subcutaneous scattering.

- (n) A green stand mixer with a white bowl attached rotates against a black background. The mixer has a sleek design with a curved neck and a whisk attachment spinning inside the bowl. The bowl is transparent, allowing a clear view of the mixing process. The mixer’s body is a soft, pastel green color, giving it a modern and clean aesthetic. The camera captures the mixer from various angles as it spins, highlighting its smooth and polished surface. Light fuzz, intact clear coat, with inner and subcutaneous scattering.
- (o) A 3D-rendered watermelon-flavored popsicle rotates against a black background. The popsicle has a glossy surface with light fuzz, intact clear coat, and subtle inner and subcutaneous scattering. It features a pink body with green rind and black seeds, giving it a realistic texture. The popsicle stick is a pale beige color, adding to the overall aesthetic. The animation showcases the popsicle from multiple angles, highlighting its smooth and shiny surface. Light fuzz, intact clear coat, with inner and subcutaneous scattering.
- (p) A cluster of purple grapes rotates against a black background, showcasing their glossy, smooth surface and light fuzz. The grapes are tightly packed together, with a few green leaves attached, adding a touch of freshness. The camera captures the grapes from various angles, highlighting their round shape and the subtle reflections. The scene is simple yet visually appealing, emphasizing the natural beauty of the fruit. Light fuzz, intact clear coat, with inner and subcutaneous scattering.

4.3 Prompt variations

Below are the full prompts for the prompt variation example (featuring the Monkey scene) in our paper.

Fuzz A 3D animated monkey character with a soft, matte finish rotates against a black background. The monkey has a round head with small ears and a tail, and it appears to be in a playful pose, moving its arms and legs in a lively manner. The animation captures the character from multiple angles as it spins, showcasing its simple yet charming design. The material has a subtle fuzziness, with inner and subcutaneous scattering, adding a touch of texture to the otherwise smooth surface. Fuzz, with inner and subcutaneous scattering.

Clearcoat A 3D animated monkey character rotates against a black background. The monkey has a round head with small ears and a tail, and it appears to be in a playful pose, moving its arms and legs in a lively manner. The animation captures the character from multiple angles as it spins, showcasing its simple yet charming design. The material has a shiny clear coat.

Dust + clearcoat A 3D animated monkey character rotates against a black background. The monkey has a round head with small ears and a tail, and it appears to be in a playful pose, moving its arms and legs in a lively manner. The animation captures the character from multiple angles as it spins, showcasing its simple yet charming design. The material has heavy dust and a clear coat with scratches. There is dust accumulated at the top of the monkey’s head.




Figure 25: Example BRDF patterns computed wave simulations. (a) A single simulated BRDF. (b) The smoother average BRDF of many similar rough surfaces. (c) The ratio between (a) and (b), which can be understood as a noise multiplier.

4.4 Microscale Enhancement Details

As mentioned in the main paper, we can additionally apply microscale procedural enhancement to materials. An example we have demonstrated is the addition of colored glints to glossy materials, which we realize using wave optics informed noise functions inspired by Yu et al. [2024, 2025]. Specifically, we can modulate the non-diffuse component of the BRDF as:

$$f_{\text{neu}}^{(i)}(\lambda, \omega_i, \omega_o) = f_{\text{neu}}(\omega_i, \omega_o) \cdot \mathcal{N}^{(i)}(\lambda, \omega_i, \omega_o) \quad (\text{S.8})$$

where $\mathcal{N}^{(i)}$ indicates we use different noise instances at different shading points. We now provide details on our noise model.

As discussed in Yu et al. [2024], BRDF patterns of rough surface samples computed from wave simulations allow for useful simplifications of the 5D noise function in Eq. S.8. In particular, the scattering value varies primarily with the half vector between ω_i and ω_o , while changing the wavelength λ mainly rescales the granular pattern. We therefore reduce the noise function to

$$\mathcal{N}^{(i)}(\lambda, \omega_i, \omega_o) := \mathcal{N}^{(i)}(\omega_x, \omega_y) \quad (\text{S.9})$$

where ω_x and ω_y are defined as

$$\omega_x = \frac{(\omega_i \cdot \mathbf{x}) + (\omega_o \cdot \mathbf{x})}{\lambda/\lambda_0}; \quad \omega_y = \frac{(\omega_i \cdot \mathbf{y}) + (\omega_o \cdot \mathbf{y})}{\lambda/\lambda_0} \quad (\text{S.10})$$

with $\lambda_0 = 400\text{nm}$ the minimum wavelength considered in our model.

To construct such a 2D noise model, we note that $\mathcal{N}^{(i)}(\omega_x, \omega_y)$ should statistically resemble the result of dividing an individual scattering pattern from wave simulation by the average over many such patterns, since it serves as a multiplier to an average BRDF. Example visualizations of a single simulated BRDF,




Figure 26: Some simplex noise tiles at different frequencies (left) and the noise pattern constructed using Eq. S.11 (right).

the smoother average BRDF of many similar rough surfaces, and their ratio are shown in Fig. 25. We seek a compact procedural model for $\mathcal{N}^{(i)}(\omega_x, \omega_y)$, and find that Simplex noise provides a lightweight construction. In particular, a weighted sum of several 2D Simplex octaves effectively reproduces the desired multiplier pattern in Fig. 25 (c). Specifically, we use

$$M^{(i)}(\omega_x, \omega_y) = \sum_{k=0}^3 \alpha^{-k} \text{snoise2} \left(\frac{\omega_x + o_{x,k}^{(i)}}{f_k}, \frac{\omega_y + o_{y,k}^{(i)}}{f_k} \right); \quad (\text{S.11})$$

$$\mathcal{N}^{(i)}(\omega_x, \omega_y) = \frac{(M^{(i)}(\omega_x, \omega_y)^2)^\beta}{\mu(\beta)}$$

Here, $\{f_k\}$ are octave frequencies that control the angular-domain correlations in noise values, and therefore in BRDF values, while α controls amplitude decay across octaves. In our example, we use geometrically spaced frequencies with ratio 1.5 and set $\alpha = 1.2$. The offsets $(o_{x,k}^{(i)}, o_{y,k}^{(i)})$ decorrelate Simplex tiles across octaves and across different noise instances. Furthermore, the exponent β controls the perceived strength of colorful glints by adjusting the variance of $\mathcal{N}^{(i)}$ across instances. The factor $\mu(\beta)$ normalizes the multiplier such that $\mathbb{E}[\mathcal{N}^{(i)}] \approx 1$ at each (ω_x, ω_y) . Numerically, we find that this factor is accurately approximated by

$$\mu(\beta) = 0.0869\beta^4 - 0.4366\beta^3 + 0.9923\beta^2 - 1.0359\beta + 0.8871. \quad (\text{S.12})$$

An illustration of our Simplex-based noise model is shown in Fig. 26. Using the 2D noise function in Eq.

S.11, we modulate the neural BRDF as

$$f_{\text{neu}}^{(i)}(\lambda, \omega_i, \omega_o) = f_{\text{neu}}(\omega_i, \omega_o) \mathcal{N}^{(i)}(\omega_x, \omega_y) \quad (\text{S.13})$$

References

- Pontus Andersson, Jim Nilsson, Tomas Akenine-Möller, Magnus Oskarsson, Kalle Åström, and Mark D. Fairchild. 2020. FLIP: A Difference Evaluator for Alternating Images. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 3, 2 (2020), 15:1–15:23. doi:10.1145/3406183
- M. Born and E. Wolf. 2013. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Pergamon. <https://books.google.ca/books?id=HY-GDAAAQBAJ>
- Ole Gulbrandsen. 2014. Artist Friendly Metallic Fresnel. *Journal of Computer Graphics Techniques (JCGT)* 3, 4 (9 December 2014), 64–72. <http://jcgt.org/published/0003/04/03/>
- Christopher Kulla and Alejandro Conty. 2017. Revisiting Physically Based Shading at Imageworks. In *SIGGRAPH 2017 Course: Physically Based Shading in Theory and Practice*.
- Christophe Schlick. 1994. An Inexpensive BRDF Model for Physically-based Rendering. *Computer Graphics Forum* 13, 3 (1994), 233–246. arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/1467-8659.1330233> doi:10.1111/1467-8659.1330233
- Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. 2007. Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques* (Grenoble, France) (*EGSR'07*). Eurographics Association, Goslar, DEU, 195–206.
- Jianfeng Xiang, Xiaoxue Chen, Sicheng Xu, Ruicheng Wang, Zelong Lv, Yu Deng, Hongyuan Zhu, Yue Dong, Hao Zhao, Nicholas Jing Yuan, and Jiaolong Yang. 2025. Native and Compact Structured Latents for 3D Generation. *Tech report* (2025).
- Yunchen Yu, Bruce Walter, Steve Marschner, and Andrea Weidlich. 2025. Realistic Cloth Rendering with a Ray-Wave Hybrid Shading Model. *ACM Transactions on Graphics (TOG)* 44, 6 (2025), 1–17.
- Yunchen Yu, Andrea Weidlich, Bruce Walter, Eugene d'Eon, and Steve Marschner. 2024. Appearance Modeling of Iridescent Feathers with Diverse Nanostructures. *ACM Transactions on Graphics (TOG)* 43, 6 (2024), 1–18.
- Tizian Zeltner, Fabrice Rousselle, Andrea Weidlich, Petrik Clarberg, Jan Novák, Benedikt Bitterli, Alex Evans, Tomáš Davidovič, Simon Kallweit, and Aaron Lefohn. 2024. Real-time Neural Appearance Models. *ACM Trans. Graph.* 43, 3, Article 33 (June 2024).