Supplementary document

May 31, 2018

Abstract

This document contains additional information to accompany paper1010: A Composite BRDF Model for Hazy Gloss.

1 BRDF Explorer shaders

We provide implementations of our composite BRDF model and of existing related models as shaders in BRDF Explorer [Dis11] in supplemental material, and encourage the reader to edit their parameters and compare their behaviors:

- `hazyGlossDielectric.brdf` implements our composite model with either Beckmann or GGX distributions, and the optional use of a coumpound masking-shadowing term. It is tailored to dielectric materials as it forbids colored peak reflectivities or edge tint, which simplifies the editing of parameters in BRDF Explorer using sliders.
- `hazyGlossConductor.brdf` implements a conductor material using colored reflectivity and edge tint parameters; the shader is otherwise identical to `hazyGlossDielectric.brdf`.
- `hazyGlossAshikhmin.brdf` implements our approach by relying on the Ashikhmin-Shirley BRDF model [AS00]; like `hazyGlossDielectric.brdf`, it is tailored to dielectric materials.
- `STD.brdf` is an implementation provided by Ribardière et al. [RBMS17] of the Cook-Torrance BRDF model using their Student-t distribution.
- `ABC.brdf` is our implementation of the microfacet-based BRDF of Löw et al. [LKYU12].
- `SGD.brdf` is our implementation of the Shifted Gamma distribution model of Bagher et al. [BSH12]. We make two simplifications to enable interactive parameter editing: we set $\Gamma = 1$ and use the Cook-Torrance masking-shadowing term. Compared to the original SGD, our version is not physically-based; however it is also simpler to edit manually and requires no precomputation. Most importantly, these simplifications do not change the outcome of our comparison in Figure 7 of the paper: when it comes to the control of haziness, the SGD model requires to manipulate two parameters, while all other models require only one.

We have supplied two live video captures of these shaders in supplemental material: one showing `hazyGlossDielectric.brdf` in action, the other showing existing BRDF models.

2 Geogebra worksheets

All the plots in the paper as well as in this document have been generated with Geogebra (`https://www.geogebra.org/`). We provide the following worksheets in supplementary material so that the interested reader might experiment with parameter editing:

- `BeckmannPhysicalBRDF.gbb` and `BeckmannPerceptualDecomp.gbb` correspond to the left and right columns of Figure 2 in the paper. We focus on the Ward BRDF model (i.e., using the Beckmann distribution) in the $\theta_d = 0$ configuration. For the sake of illustration, we replace $4 \cos\theta_i \cos\theta_o$ in the denominator of each component BRDF by 4 and their masking-shadowing term G by 1. The resulting simplified BRDF is a function of θ_h only.
- `haloParams.gbb` corresponds to Figure 3 in the paper. It plots the haze intensity at normal incidence k_h as a function of the core reflectivity r_c, shown on the horizontal axis. Both the piecewise linear and smooth interpolation are plotted, and controlled by the peak ratio p, haziness β_h and smoothness w.

- `peakFunc.gbb` corresponds to Figure 1 in Section 3 of this document. It shows the peak removal function P with respect to θ_d, along with the ratio function R.

We have supplied two live video captures showing the `BeckmannPhysicalBRDF.gbb` and `BeckmannPerceptualDecomp.gbb` worksheets in action in supplemental material.

3 Peak removal function

Our perceptual decomposition (Equation 2 in the paper) relies on a peak removal function P, which varies with θ_d. In practice, our haze mapping only requires to consider the peak ratio $p = P(\theta_d = 0)$; we explain in Section 3.1 how to obtain the p formula for common BRDF models as presented in Table 2 of the paper. We also provide an analysis of the full peak removal function P in Section 3.2.

3.1 p formula

Equation 1 in the paper defines the composite BRDF model as $f_r = ((1-\beta)f_n + \beta f_w)F_r$, where F_r is the Fresnel term, while n and w subscripts stand for 'narrow' and 'wide'. In the context of microfacet theory, the BRDF functions f_n and f_w are given by:

$$f_{n,w}(\omega_i, \omega_o) = \frac{D_{n,w}(\mathbf{h})G_{n,w}(\omega_i, \omega_o)}{4|\omega_i \cdot \mathbf{n}||\omega_o \cdot \mathbf{n}|}, \tag{1}$$

where $\mathbf{h} = \frac{\omega_i+\omega_o}{\|\omega_i+\omega_o\|}$ is the half-way vector, D_n and D_w are distribution terms, and G_n and G_w are masking-shadowing terms. Re-expressing the peak ratio $p = \frac{f_w(\mathbf{n},\mathbf{n})}{f_n(\mathbf{n},\mathbf{n})}$ in this context, we obtain:

$$p = \frac{D_w(\mathbf{n})G_w(\mathbf{n},\mathbf{n})}{D_n(\mathbf{n})G_n(\mathbf{n},\mathbf{n})} = \frac{D_w(\mathbf{n})}{D_n(\mathbf{n})}, \tag{2}$$

since $G_n(\mathbf{n},\mathbf{n}) = G_w(\mathbf{n},\mathbf{n}) = 1$ (all the microfacets are visible at normal incidence).

The exact formula for p then depends on the choice of distribution. Both the Cook-Torrance [CT82] and Ward [War92] models are based on the Beckmann distribution, the only difference being that the former is isotropic and the latter anisotropic. Since $D^{\text{Beckmann}}(\mathbf{n}) = \frac{1}{\pi\alpha^x\alpha^y\cos^4\theta_h}$, we obtain $p = \frac{\alpha_n^x\alpha_n^y}{\alpha_w^x\alpha_w^y}$ in the anisotropic case, and $p = \frac{\alpha_n^2}{\alpha_w^2}$ in the isotropic case. The model of Walter et al. [WMLT07] is based on the GGX distribution. Interestingly, $D^{\text{GGX}}(\mathbf{n}) = D^{\text{Beckmann}}(\mathbf{n})$, hence once again $p = \frac{\alpha_n^x\alpha_n^y}{\alpha_w^x\alpha_w^y}$ in the anisotropic case, and $p = \frac{\alpha_n^2}{\alpha_w^2}$ in the isotropic case.

The Blinn-Phong BRDF model [Bli77] rather makes use of a shininess parameter s that controls the sharpness of highlights. Its normalization factor is given by $D^{\text{Blinn-Phong}}(\mathbf{n}) = \frac{(s+2)(s+4)}{8\pi(2^{-s/2}+s)}$, yielding $p = \frac{(s_w+2)(s_w+4)(2^{-s_n/2}+s_n)}{(s_n+2)(s_n+4)(2^{-s_w/2}+s_w)}$. The Ashikhmin-Shirley BRDF model [AS00] generalizes Phong's to the anisotropic case with $D^{\text{Ashikhmin-Shirley}}(\mathbf{n}) = \frac{\sqrt{(s^x+1)(s^y+1)}}{8\pi}$, yielding $p = \frac{\sqrt{(s_w^x+1)(s_w^y+1)}}{\sqrt{(s_n^x+1)(s_n^y+1)}}$.

3.2 Peak analysis

The full peak removal function P is a function of θ_d; from Equation 3 in the paper, we obtain:

$$P(\theta_d) = \frac{F_c(\theta_d) - (1-\beta)F_r(\theta_d)}{\beta F_r(\theta_d)}. \tag{3}$$

As shown in the top left of Figure 1, P departs from p with increasing θ_d until it reaches $P(\frac{\pi}{2}) = 1$. As P departs from p, the specular core and surrounding halo are not perfectly independently controllable anymore, but the effect of one on the other remains negligible in practice.

The effect on the decomposition may be illustrated by the ratio of the BRDF to the specular core *in the specular direction*. We denote this ratio by $\mathrm{R} = \frac{f(\mathbf{n},\mathbf{d})}{f_n(\mathbf{n},\mathbf{d})F_c(\theta_d)}$, where both f_n and f_w are expressed in the halfway-difference parametrization $(\mathbf{h},\mathbf{d})$, with $\mathbf{h} = \mathbf{n}$ since we only consider the specular direction. We then obtain:

$$\mathrm{R}(\theta_d) = 1 + \frac{K_h(\theta_d)}{F_c(\theta_d)}\left(\frac{f_w(\mathbf{n},\mathbf{d})}{f_n(\mathbf{n},\mathbf{d})} - P(\theta_d)\right). \tag{4}$$

For most BRDF models, we have $\frac{f_w(\mathbf{n},\mathbf{d})}{f_n(\mathbf{n},\mathbf{d})} = p$ since the masking-shadowing term is independent of roughness parameters, and hence vanishes in the ratio. An important exception is

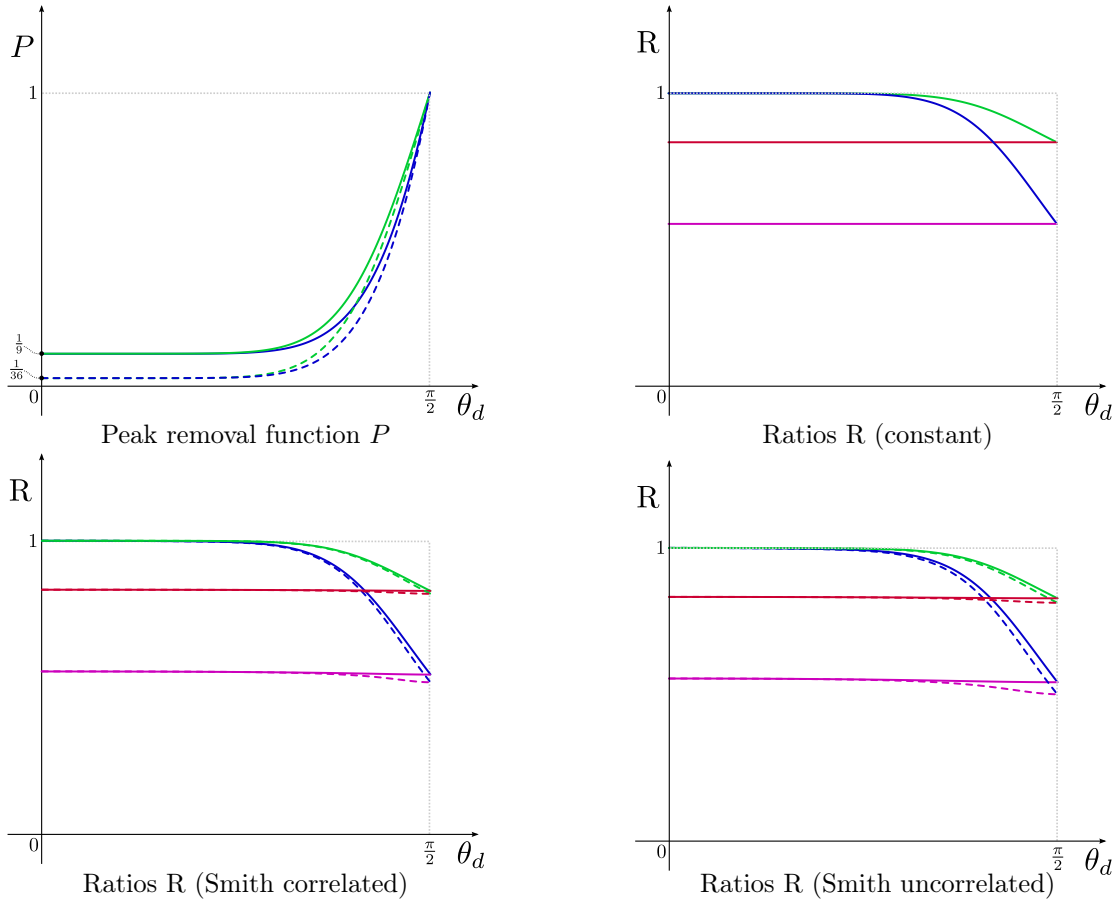

Figure 1: The top left plot shows the peak removal function P as a function of θ_d. The solid curves correspond to $\lambda_h = 2$ and the dashed curves to $\lambda_h = 5$, yielding $p = \frac{1}{9}$ and $p = \frac{1}{36}$ respectively. Green curves are obtained by setting $\beta_h = 0.8$; blue curves are obtained with $\beta_h = 0.2$. The other parameters remain fixed at $r_c = 0.5$ and $\alpha_n = 0.1$. The three remaining figures show the ratio R for three types of shadowing-masking terms. There are no dashed curves in the top right figure since R is independent of roughness in the constant case. The variations of R with λ_h (solid vs dashed lines) in the other cases are subtle. We also plot $\frac{f_r}{f_n}$ using red and pink curves for reference. Observe that green-red and blue-pink pairs of curves meet at $\theta_d = \frac{\pi}{2}$ as expected.

Smith's masking-shadowing model, which does depend on roughness through the function Λ (see [Hei14]). In our case, Λ is parametrized by θ_d since we only consider the $\theta_h = 0$ configuration; hence both $\frac{f_w}{f_n}$ and R are indeed functions of θ_d. Smith's model may be defined either considering that masking and shadowing are independent, or that they are height-correlated. In its uncorrelated form, $\frac{f_w}{f_n}(\theta_d) = p\left(\frac{1+\Lambda_n(\theta_d)}{1+\Lambda_w(\theta_d)}\right)^2$. In its correlated form, $\frac{f_w}{f_n}(\theta_d) = p\,\frac{1+2\Lambda_n(\theta_d)}{1+2\Lambda_w(\theta_d)}$. In both cases, Λ_n and Λ_w correspond to integrals over slopes of the narrow and wide distributions respectively. The case of a compound masking-shadowing term as described in Section 3.2 of the paper is actually simpler, since the common masking-shadowing term vanishes in the ratio, yielding $\frac{f_w(\mathbf{n},\mathbf{d})}{f_n(\mathbf{n},\mathbf{d})} = p$. Figure 1b-d visualizes R (blue and green curves) for the three types of masking-shadowing terms: independent of roughness, Smith height-correlated and Smith uncorrelated.

Three observations may be made from these figures. First, the differences between various types of masking-shadowing terms are not significant. Second, R is a monotonically decreasing function of θ_d: when $\theta_d > 0$, the peak value is reduced by the addition of a halo. Third, the worst separation between the specular core and halo occurs when $\theta_d = \frac{\pi}{2}$:

$$\mathrm{R}(\tfrac{\pi}{2}) = 1 + \beta\left(\tfrac{f_w}{f_n}(\tfrac{\pi}{2}) - 1\right) = \frac{f_r}{f_n}(\tfrac{\pi}{2}). \tag{5}$$

R then becomes equivalent to the ratio of the BRDF to the narrow component in the specular direction at $\theta_d = \frac{\pi}{2}$. We visualize this ratio with red and pink curves in Figure 1b-d. In other words, we revert to the case of the physically-based decomposition at extreme grazing angles.

The extent of the plateau in R (i.e., the range of θ_d values where the central peak and halo are quasi indepentent) is determined by the choice of material parameters. The interested reader is encouraged to edit these parameters using the `peakFunc.gbb` worksheet.

4 Relationship to Vangorp et al.

The following derivation is provided for comparison to the work of Vangorp et al. [VBF17]. They define a halo energy $E_h = 1 - \frac{r_c}{r}$, which they show is correlated to their experimental perceptual haziness ratings.

The halo energy has two drawbacks when it comes to artistic material editing. First, if we rewrite the positivity and energy conservation constraints in terms of E_h, we obtain $E_h \leq 1 - p$ and $E_h \leq 1 - r_c$ respectively. This means that the range of values the halo energy can take must be limited in a way that depends either on roughnesses (through p) or on the core reflectivity r_c. Second, since we have $r = \frac{r_c}{1-E_h}$, uniform variations in the halo energy E_h will yield non-uniform variations of the reflectivity r; in particular, it will increase suddenly for higher E_h values.

Using our haziness parameter β_h, the halo energy may be rewritten as:

$$E_h = \begin{cases} \frac{\beta_h(1-p)}{p+\beta_h(1-p)} & \text{if } r_c \leq p \\ \frac{\beta_h(1-r_c)}{r_c+\beta_h(1-r_c)} & \text{otherwise.} \end{cases}$$

Intuitively, the haziness $\beta_h \in [0, 1]$ is first linearly remapped to either $[0, 1 - p)$ (`P` constraint) or $[0, 1 - r_c)$ (`E` constraint); then a non-linearity is applied to counteract the non-uniform variations of r when E_h is modified. This demonstrates the advantage of using β_h instead of E_h to control the halo component: it provides a linear and well-bounded behavior.

5 Additional rendering results

5.1 Spatially-varying core reflectivity

Figures 2 and 3 compare the two strategies for computing k_h, the haze intensity at $\theta_d = 0$. Using a piecewize formula (Equation 7 in the paper) produces a visual discontinuity when the core reflectivity r_c varies spatially. The discontinuity is smoothed out by using rational Bézier interpolation (Equation 15 in the appendix of the paper), which also tends to reduce the haze intensity k_h. Note however that from a physical point of view, such spatial variations of r_c are uncommon since they imply a spatially-varying refractive index.

5.2 Compound masking-shadowing term

We show in Figure 4 (for the GGX distribution) and 5 (for the Beckmann distribution) the effect of choosing a compound masking-shadowing term over separate terms for each component (we use $r_c = 0.5$, $\alpha_n = 0.03$, $\beta_h = 0.99$, $\lambda_h = 10$ in both cases). As discussed in Section 3 of the paper, a compound term corresponds to intertwined narrow and wide microfacet distributions, while separate terms corresponds to having large patches for each distribution. The visual effect of the compound term is increased intensity at grazing angles, in particular with the GGX distribution.

5.3 Anisotropic hazy gloss

Figure 6 is an expanded version of Figure 4 in the paper. It illustrates the versatility of our approach when it comes to anisotropic materials: either the specular core and/or surrounding halo may be made isotropic or anisotropic, providing interesting artistic controls over hazy gloss.

5.4 Global illumination results

In Figure 7, we present the same results as those shown in Figure 1 in the paper with the addition of a colored conductor material, this time with full-size images for the hazy and haze-free versions. We show additional rendering results using a transparent material in Figure 8.

5.5 Haziness editing at the compositing stage

Using the approximation of Section 4.3 in the paper, we may grant editing of haziness at the compositing stage. This is shown in Figure 9 (focusing on the seat cushion): four render buffers are used for the specular component, plus another buffer for the diffuse component. There are then recombined according to Equation 11 in the paper with arbitrary core reflectivity and haziness parameters. The inset images show differences with ground renderings computed using the exact Fresnel equations (as opposed to our approach that uses Schlick's approximation). They have been computed in HDR and their absolute value multiplied by 100, showing that differences between compositing and rendering on the cushion are negligible. Our simple compositing method does not affects inter-reflections on the interior of the seat, but this is clearly out of the scope of this work.

5.6 Three-component specular BRDF model

To evaluate the visual impact of adding a third component to Equation 1 in the paper, we rewrite it as $f_r = \big((1-\beta-\beta_2)f_n + \beta f_w + \beta_2 f_{w_2}\big)F_r$, where f_{w_2} has a wider extent than f_w. As shown in Figure 10, the resulting material still appears to be composed of a specular core and a surrounding halo (not of three visual components) irrespective of the choice of blending parameter β_2 or roughness α_{w_2}. This suggests that a pair of components is enough for the artistic control of hazy gloss; however, additional perceptual experiments are required to validate this claim.

References

[AS00] Ashikhmin M., Shirley P.: An anisotropic Phong BRDF model. *Journal of Graphics Tools 5*, 2 (2000), 25–32.

[Bli77] Blinn J. F.: Models of light reflection for computer synthesized pictures. *SIGGRAPH Comput. Graph. 11*, 2 (July 1977), 192–198. URL: http://doi.acm.org/10.1145/965141.563893, doi:10.1145/965141.563893.

[BSH12] Bagher M. M., Soler C., Holzschuch N.: Accurate fitting of measured reflectances using a Shifted Gamma micro-facet distribution. *Computer Graphics Forum 31*, 4 (2012), 1509–1518. doi:10.1111/j.1467-8659.2012.03147.x.

[CT82] Cook R. L., Torrance K. E.: A reflectance model for computer graphics. *ACM Transactions on Graphics 1*, 1 (1982), 7–24. doi:10.1145/357290.357293.

[Dis11] Disney: Brdf explorer. https://www.disneyanimation.com/technology/brdf.html, 2011.

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

[LKYU12] Löw J., Kronander J., Ynnerman A., Unger J.: BRDF models for accurate and efficient rendering of glossy surfaces. *ACM Trans. Graph. 31*, 1 (Feb. 2012), 9:1–9:14. URL: http://doi.acm.org/10.1145/2077341.2077350, doi:10.1145/2077341.2077350.

[RBMS17] Ribardière M., Bringier B., Meneveaux D., Simonot L.: STD: Student's t-Distribution of Slopes for Microfacet Based BSDFs. *Computer Graphics Forum 36*, 2 (2017), 421–429. doi:10.1111/cgf.13137.

[VBF17] Vangorp P., Barla P., Fleming R. W.: The perception of hazy gloss. *Journal of Vision 17*, 5 (2017), 19:1–17. doi:10.1167/17.5.19.

[War92] Ward G. J.: Measuring and modeling anisotropic reflection. *Computer Graphics (Proceedings of ACM SIGGRAPH 92) 26*, 2 (1992), 265–272. doi:10.1145/142920.134078.

[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* (2007), EGSR'07, Eurographics Association, pp. 195–206. URL: http://dx.doi.org/10.2312/EGWR/EGSR07/195-206, doi:10.2312/EGWR/EGSR07/195-206.

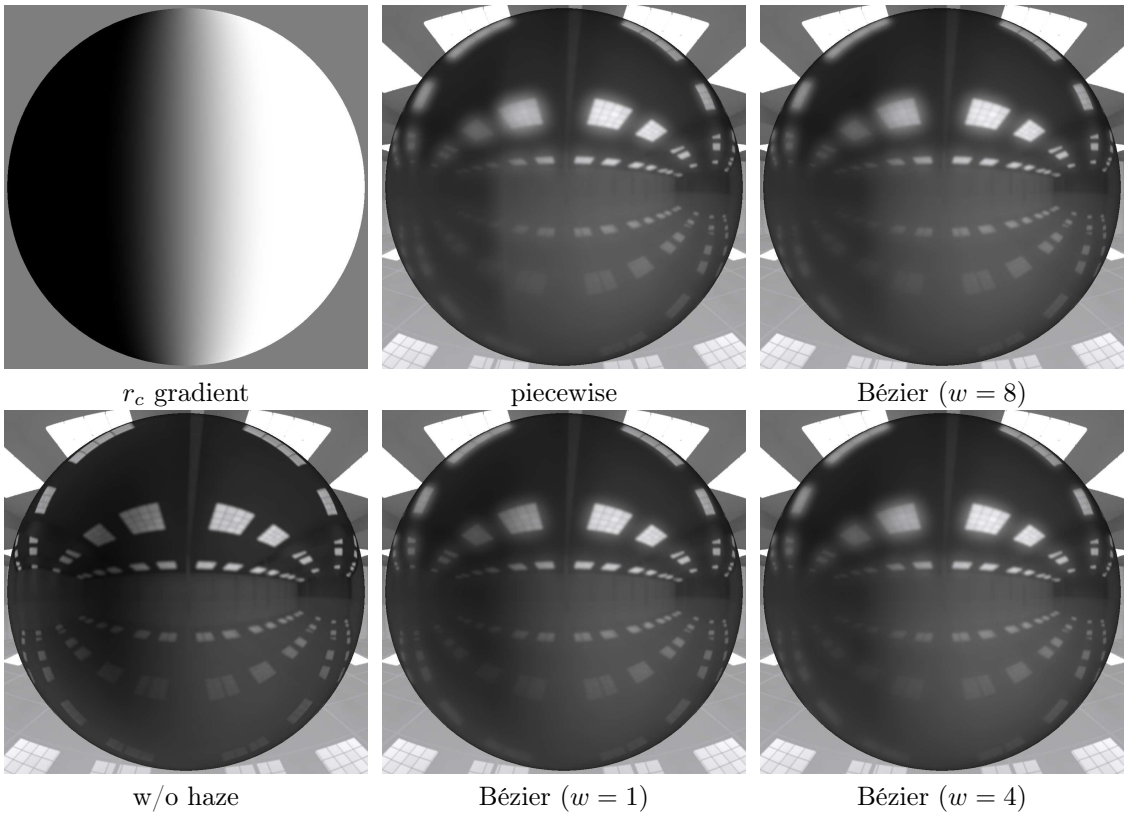

Figure 2: When the core reflectivity r_c is varied spatially (top left), the piecewise formula for computing haze intensity produces a visual discontinuity. Using rational Bézier interpolation permits to smooth it out, but with stronger smoothing (smaller w) haziness is also reduced.

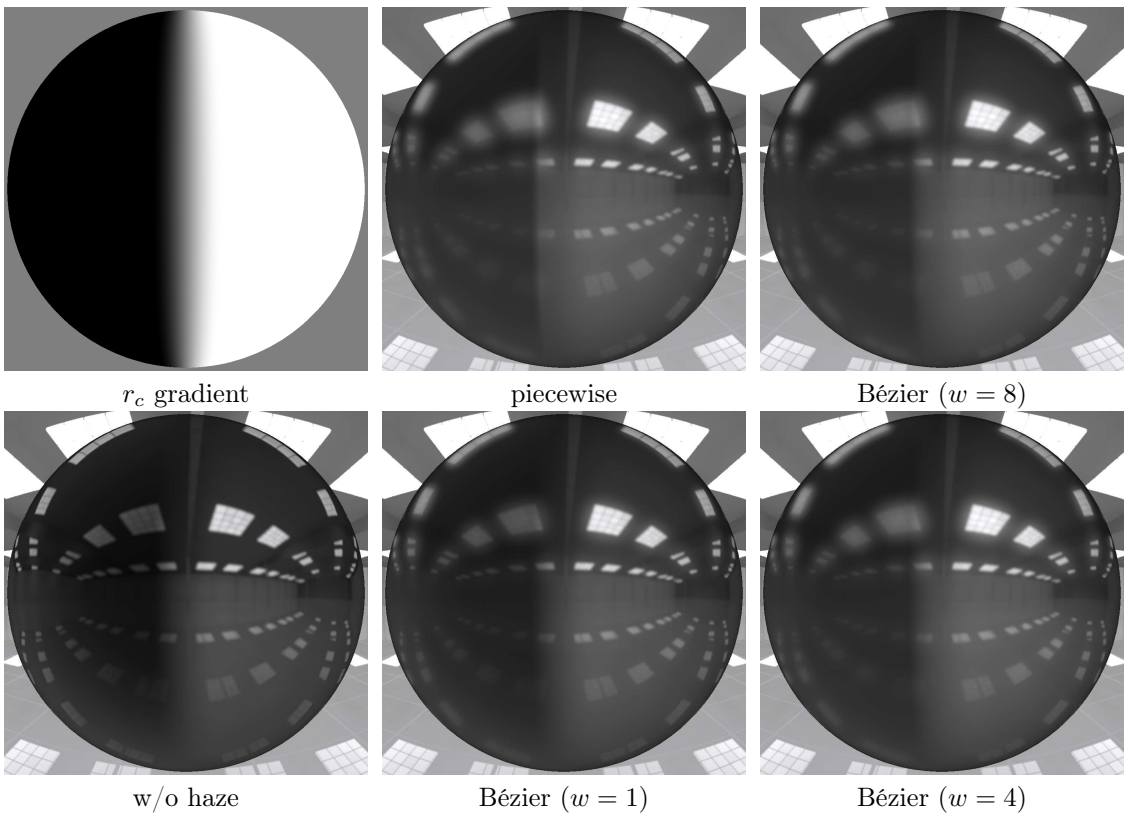

Figure 3: Same as Figure 2 but with a sharper gradient of core reflectivity r_c.

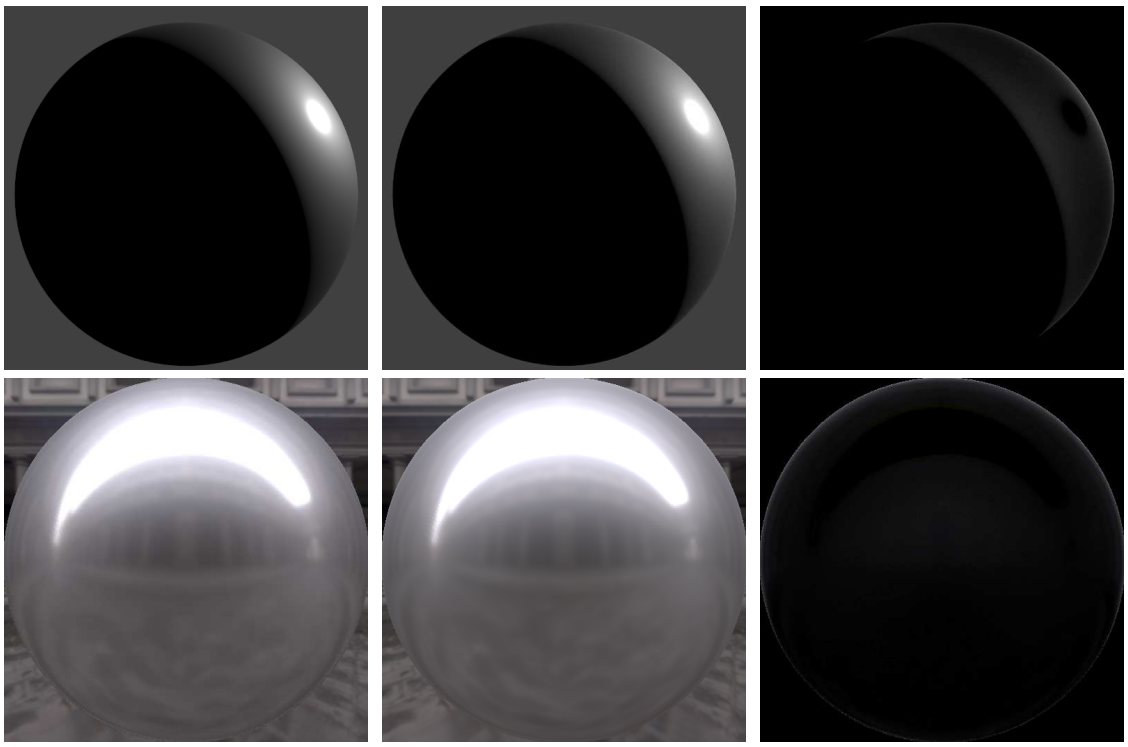

Separate masking-shadowing | Compound masking-shadowing | Difference image (×4)

Figure 4: Using the compound masking-shadowing term with the GGX distribution and a wide halo component slightly increases intensity at grazing lighting and viewing angles, both when using a single light source (top row) or the `Uffizi` environment lighting (bottom row) for rendering.

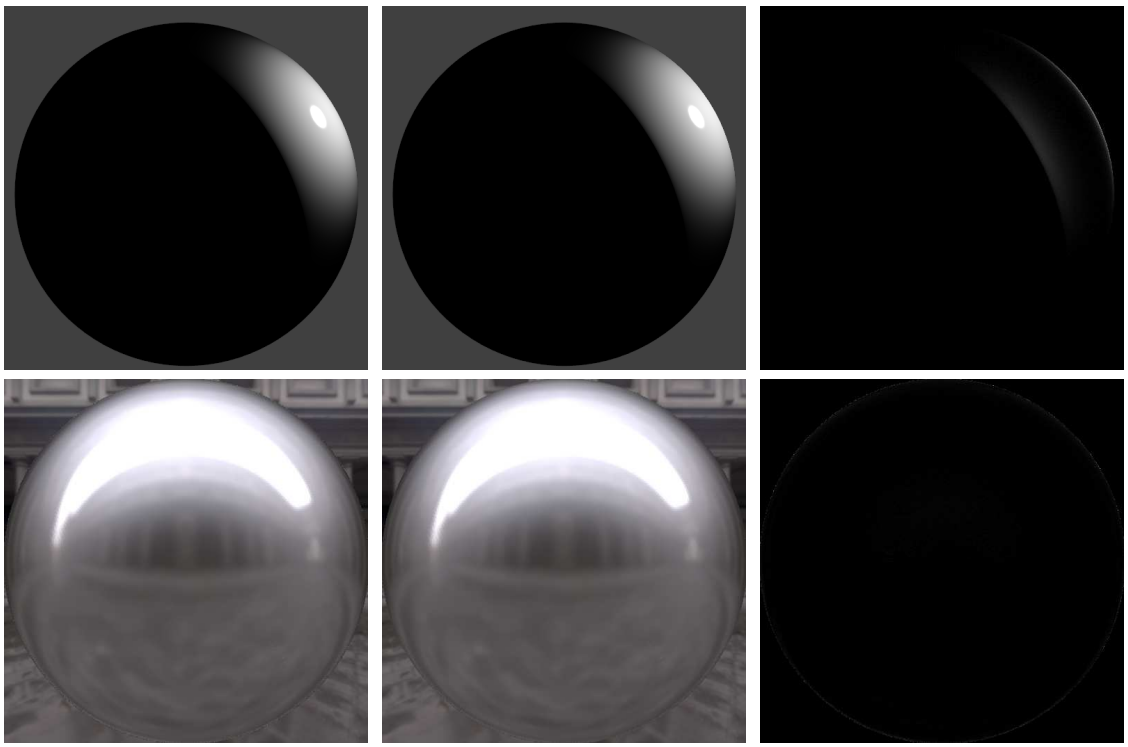

Separate masking-shadowing | Compound masking-shadowing | Difference image (×4)

Figure 5: Using the compound masking-shadowing term with the Beckmann distribution only very slightly increases intensity at grazing lighting and viewing angles when using a single light source (top row), and has a negligible effect with the chosen environment lighting (bottom row).

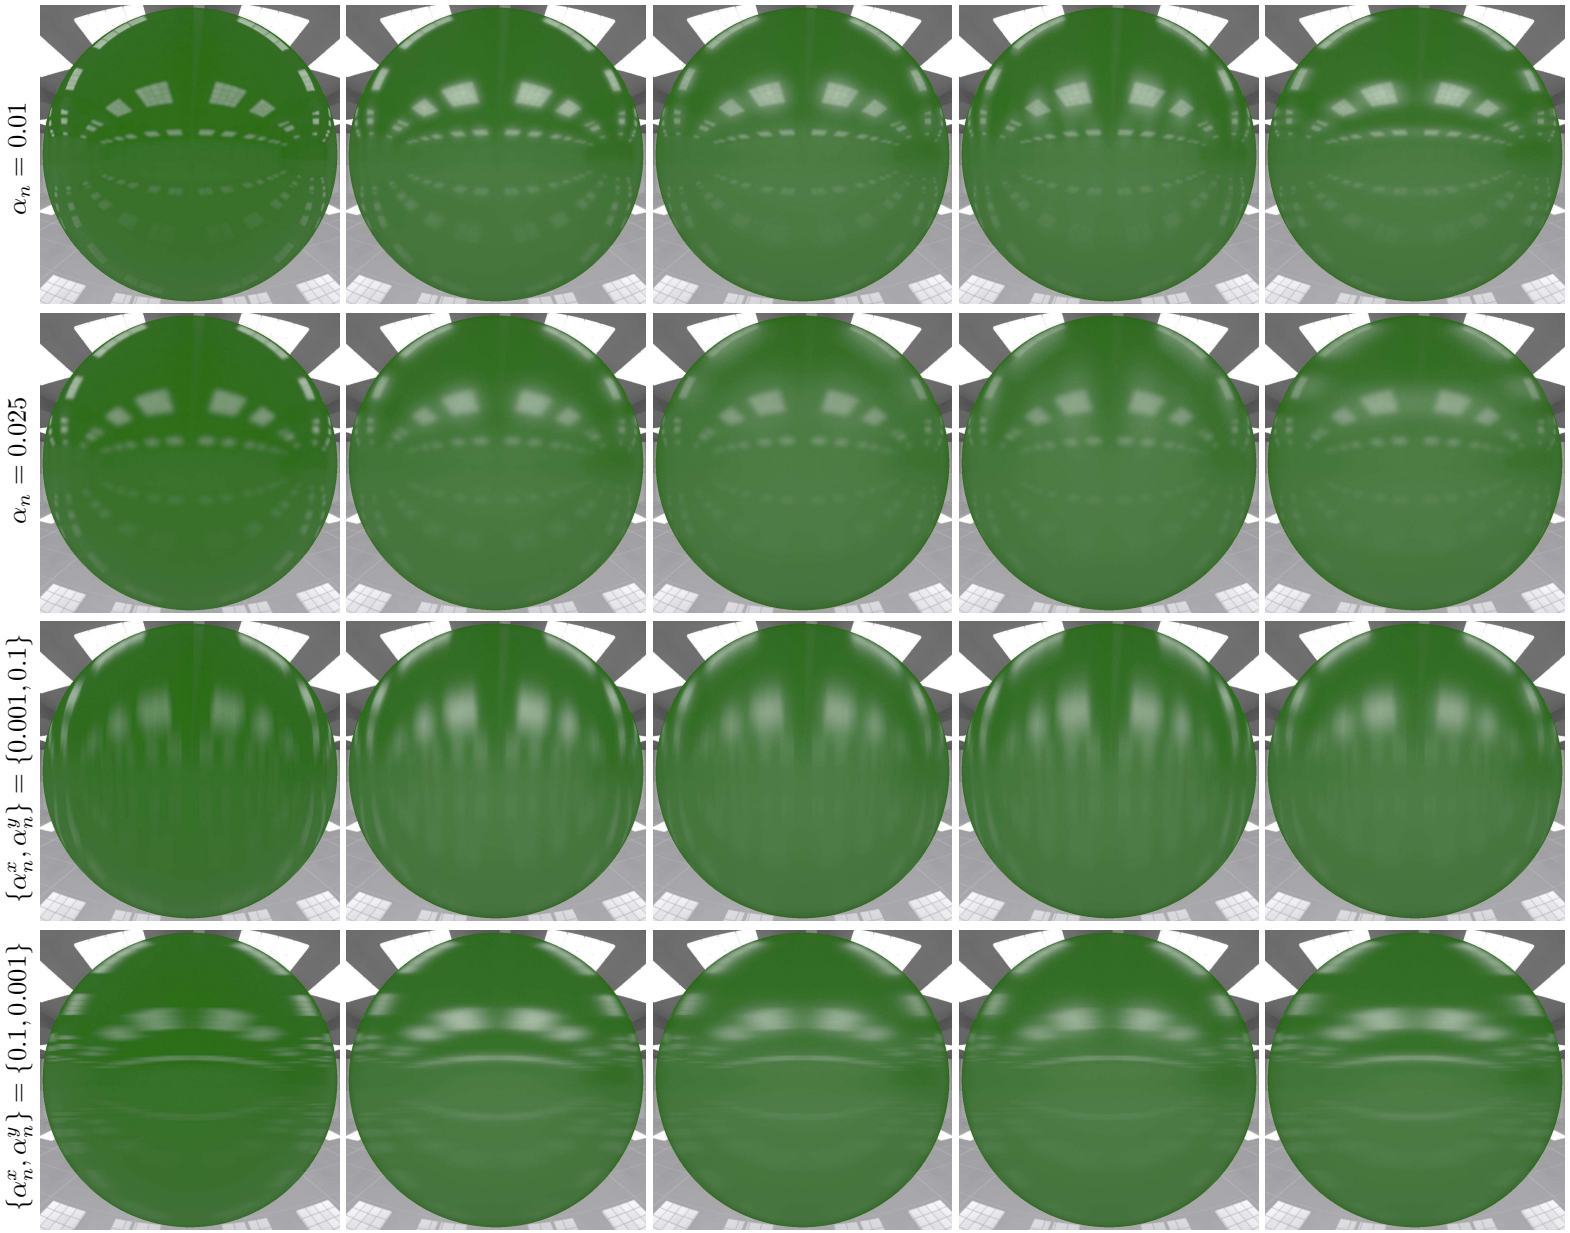

Figure 6: Various combinations of $\alpha_n^{\{x,y\}}$ and $\lambda_h^{\{x,y\}}$ for a *dielectric* material (we fix $r_c = 0.03$ and $\beta_h = 0.1$). Please zoom in to see haze effects.

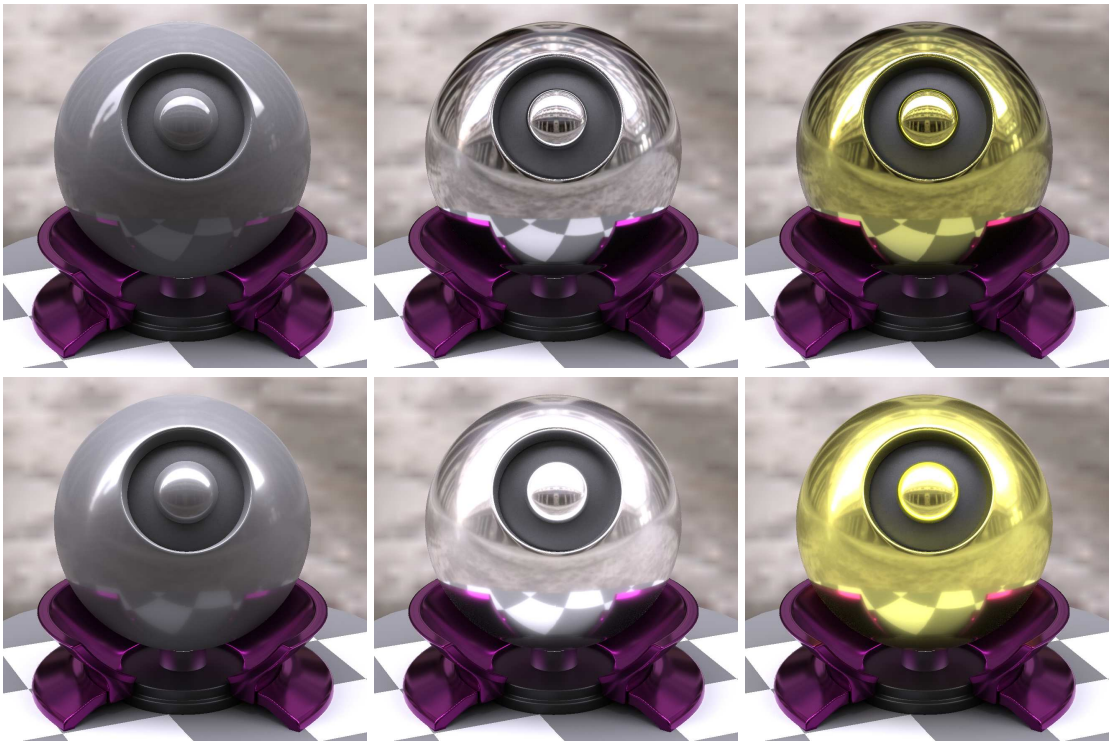

Figure 7: Renderings of the `Probe` materials: dielectric, conductor and achromatic conductor; without haze (top row) and with haze (bottom row).

Figure 8: Renderings of a transparent dielectric material ($r_c = 0.04, with \alpha_n = 0.02$ on top and $\alpha_n = 0.05$ at bottom) without haze ($\beta_h = 0$) at left, and with $\beta_h = 0.1$ in the middle (with $\lambda_h = 3$) and at right (with $\lambda_h = 5$).

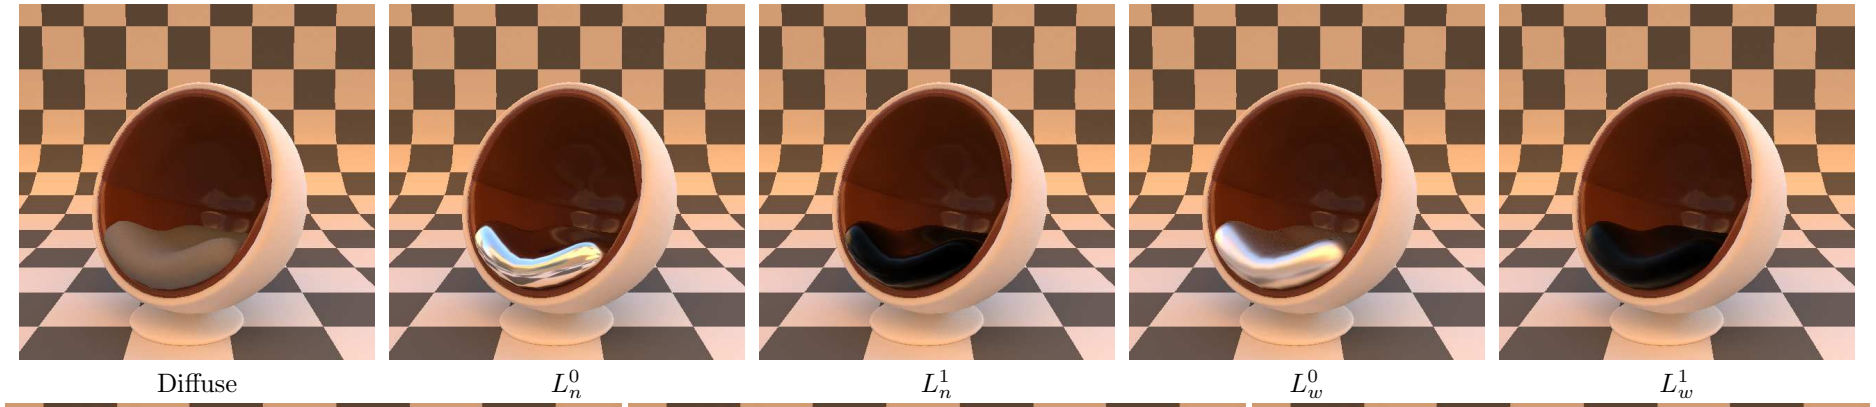

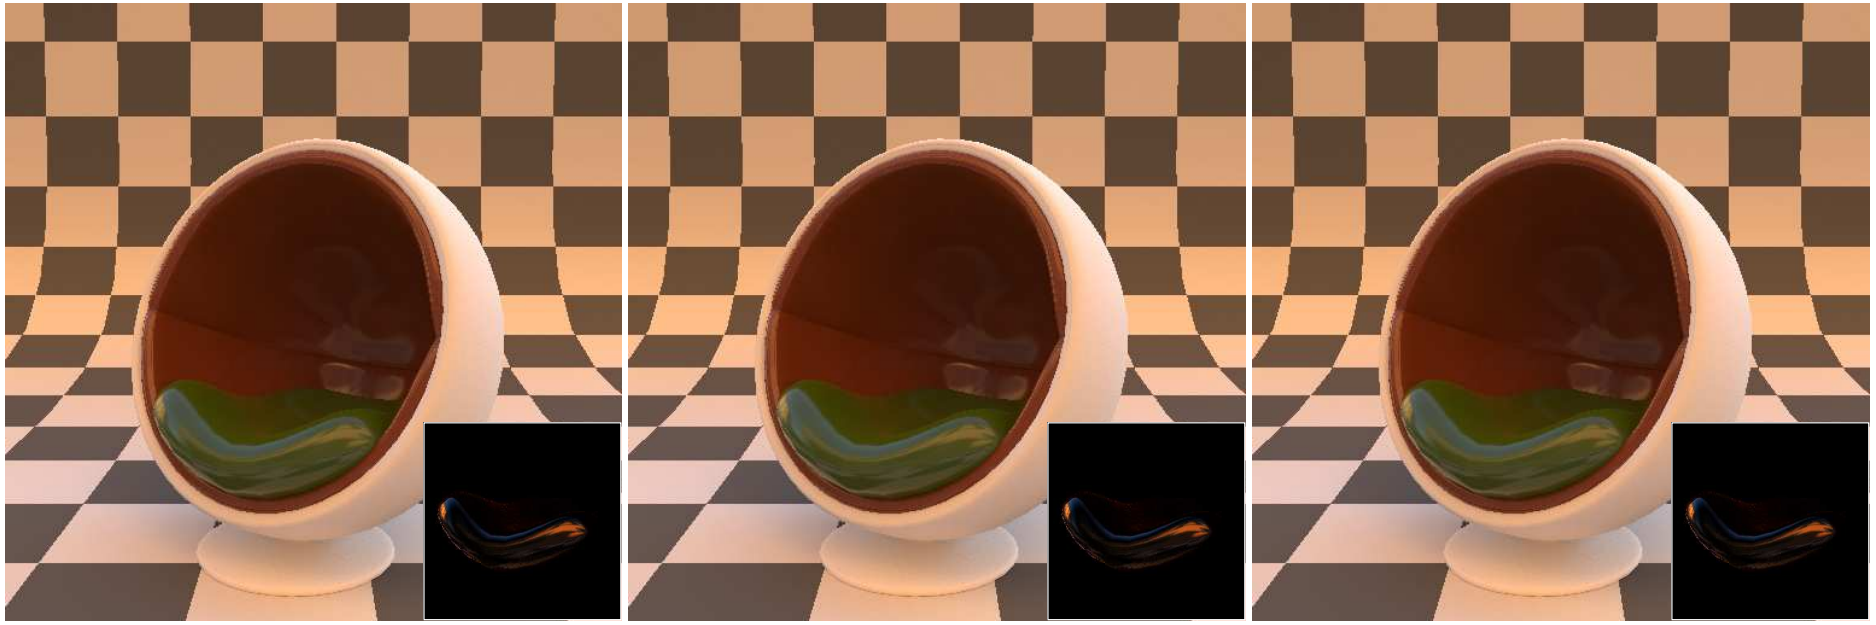

Figure 9: Top row: we take five render buffers (focusing on the seat cushion) as input to our compositing application. Bottom row: by recombining the buffers we obtain a composite image that may be edited to manipulate the haziness of one object. The process is demonstrated in the supplemental video. The insets show difference images on the cushion between our compositing results and ground truth renderings, in high dynamic range and **multiplied by 100**.

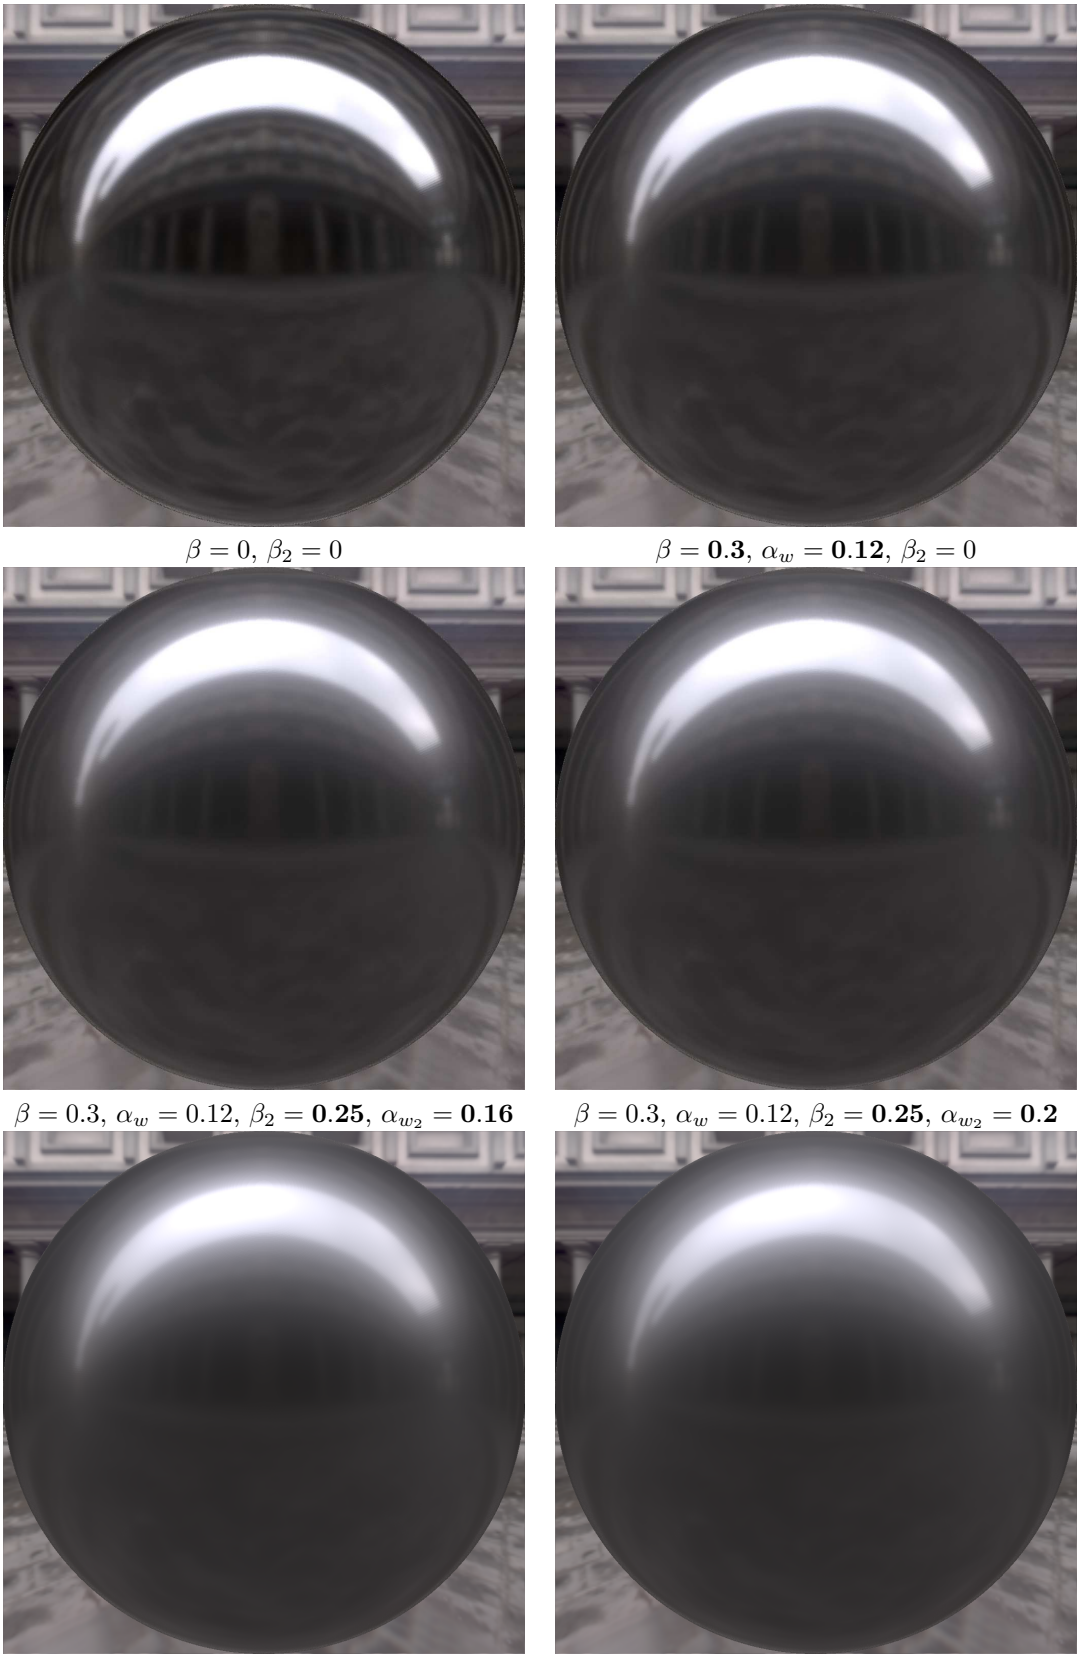

$\beta = 0$, $\beta_2 = 0$

$\beta = \mathbf{0.3}$, $\alpha_w = \mathbf{0.12}$, $\beta_2 = 0$

$\beta = 0.3$, $\alpha_w = 0.12$, $\beta_2 = \mathbf{0.25}$, $\alpha_{w_2} = \mathbf{0.16}$

$\beta = 0.3$, $\alpha_w = 0.12$, $\beta_2 = \mathbf{0.25}$, $\alpha_{w_2} = \mathbf{0.2}$

$\beta = 0.3$, $\alpha_w = 0.12$, $\beta_2 = \mathbf{0.5}$, $\alpha_{w_2} = \mathbf{0.16}$

$\beta = 0.3$, $\alpha_w = 0.12$, $\beta_2 = \mathbf{0.5}$, $\alpha_{w_2} = \mathbf{0.2}$

Figure 10: Preliminary experiments with a three-component BRDF. The first row shows a BRDF with one (left) and two (right) components. The middle and bottom rows show a three-component BRDF with different sets of parameters. In these cases, the presence of a third component does not seem to yield a visual appearance that is significantly different from hazy gloss.