


Supplemental Material: Spectral Upsampling Approaches for RGB Illumination

G. C. Guarnera^{1,2,3} , Y. Gitlina⁴, V. Deschaintre^{4,5} and A. Ghosh^{1,4}

¹Lumirithmic Ltd

²University of York, United Kingdom

³Norwegian University of Science and Technology, Norway

⁴Imperial College London, United Kingdom

⁵Adobe Research, United Kingdom

1. Additional results

We provide additional figures, not included in the main paper. We visualize results of spectral upsampling using our method on three more environment maps (Fig. 1) and present additional lighting reproduction examples in a multispectral LED sphere, comparing RGB and spectral illumination reproduced with our method (Figs. 2 and 3). We also plot RGBACAW basis employed on USC ICT's Lightstage X (Fig. 4) and show more results for our method using our RGBW+ basis (Fig. 5). We report in detail the optimal Gaussian bases as well as corresponding spectral and color checker reconstruction using these bases in our optimization procedure (Figs. 6 and 7) and provide additional comparisons to existing spectral upsampling methods for an outdoor illuminant and a reflectance spectrum (Fig. 8). We show in Table 1 the average ΔE_{94} for each basis on color checker illumination.

Fig. 9 highlights with an example the non-convexity of the optimization problem we address in our work, while Table 2 reports weights and XYZ coordinates of the spectra used in the example; Fig. 10 demonstrates that these spectra, and the corresponding weights, actually represents local minima for the ΔE_{94} objective function. Fig. 11 compares our proposed genetic algorithm based optimization with standard convex non-linear optimization for spectral upsampling and shows genetic algorithm to converge to spectra that is closer to the ground truth.

Finally, Fig. 12 shows another example of application of our RGB-to-spectral upsampling using the proposed Gaussian basis for improved color accuracy of renderings.

Basis	Average ΔE_{94}
RGBW+ (ours)	1.8658
RGBACW	2.7350
Gaussian (3-3)	1.3970
Gaussian (5-1)	1.8545

Table 1: Average ΔE_{94} (lower is better) for the different evaluated basis. We evaluate each basis on different illuminants reconstruction and its impact on color checkers. As expected, the theoretical Gaussian bases provide a bit better results than the existing physical ones, but our basis average error still remains below $\Delta E_{94} < 2$.

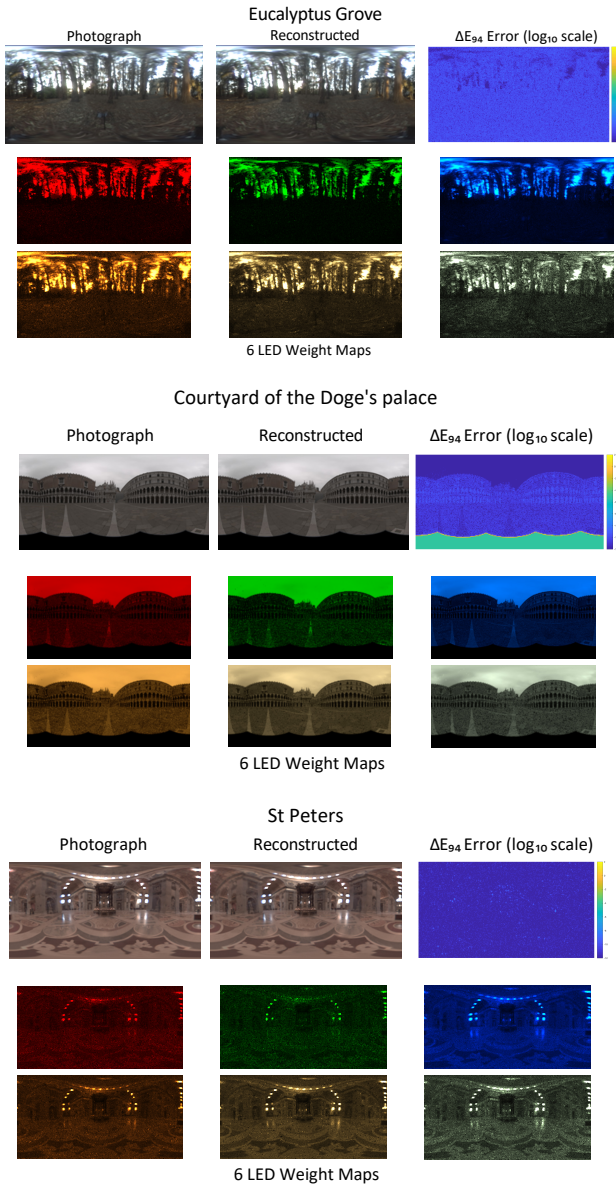


Figure 1: Spectral upsampling of additional environment maps using our method with RGBW+ basis.




Figure 2: Additional comparisons between RGB and spectral lighting reproduction with our approach on a set of color crayons, using a multispectral LED sphere. The top-row presents mirror ball photographs of the target environmental illumination. The spectral lighting reproduction (bottom row) is a closer match to the reference photograph than RGB lighting reproduction, while still being driven by RGB lightprobe as input.




Figure 3: Additional examples of lighting reproduction on a white deer figurine for both indoor and outdoor lighting environments. Our proposed spectral upsampling driven lighting reproduction is qualitatively a closer match to the reference photograph than RGB lighting reproduction.

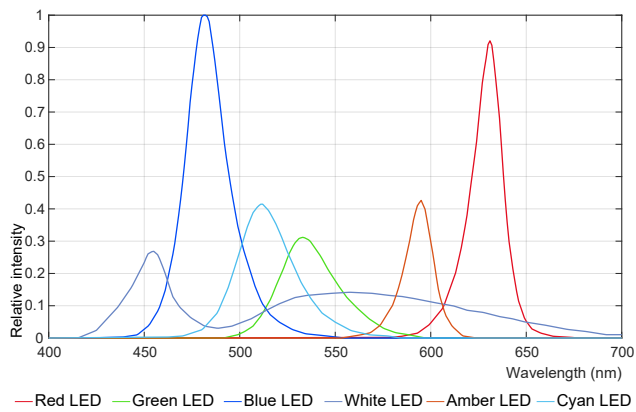


Figure 4: RGBCAW emission spectra employed on USC ICT's Lightstage X for spectral lighting reproduction. They use 5 narrowband and 1 broadband illuminants.

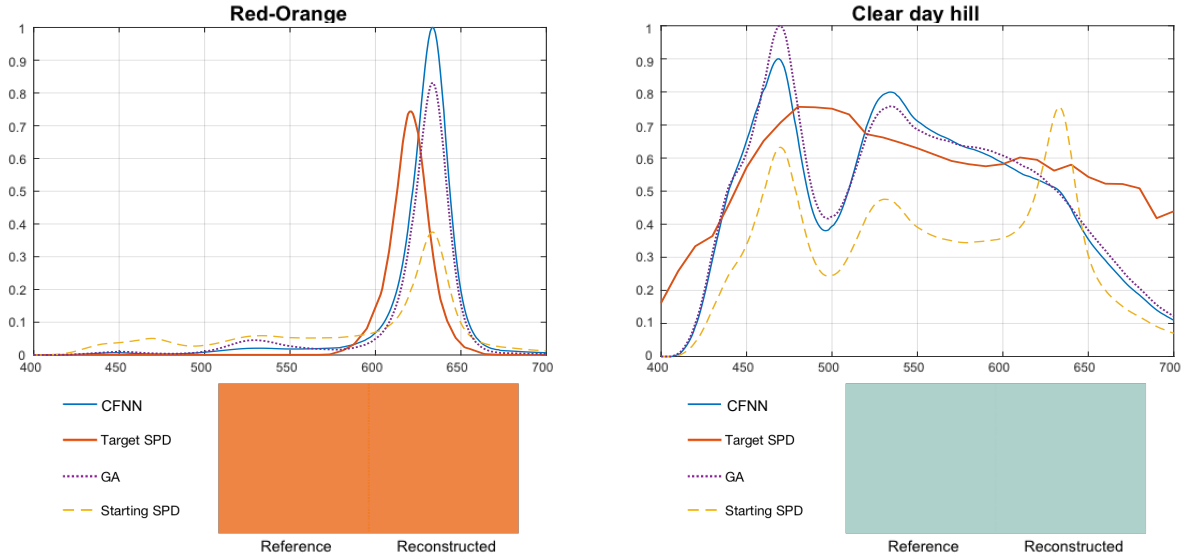


Figure 5: Additional examples of spectral and color reconstruction with our method using our RGBW+ basis.




Figure 6: Initial set of Gaussians used to determine optimal bases with different number of narrow and broad bands (left). At the center, the optimal basis consisting of 3 broad and 3 narrow bands, whereas on the right we report the optimal basis consisting of 1 broad and 5 narrow bands.

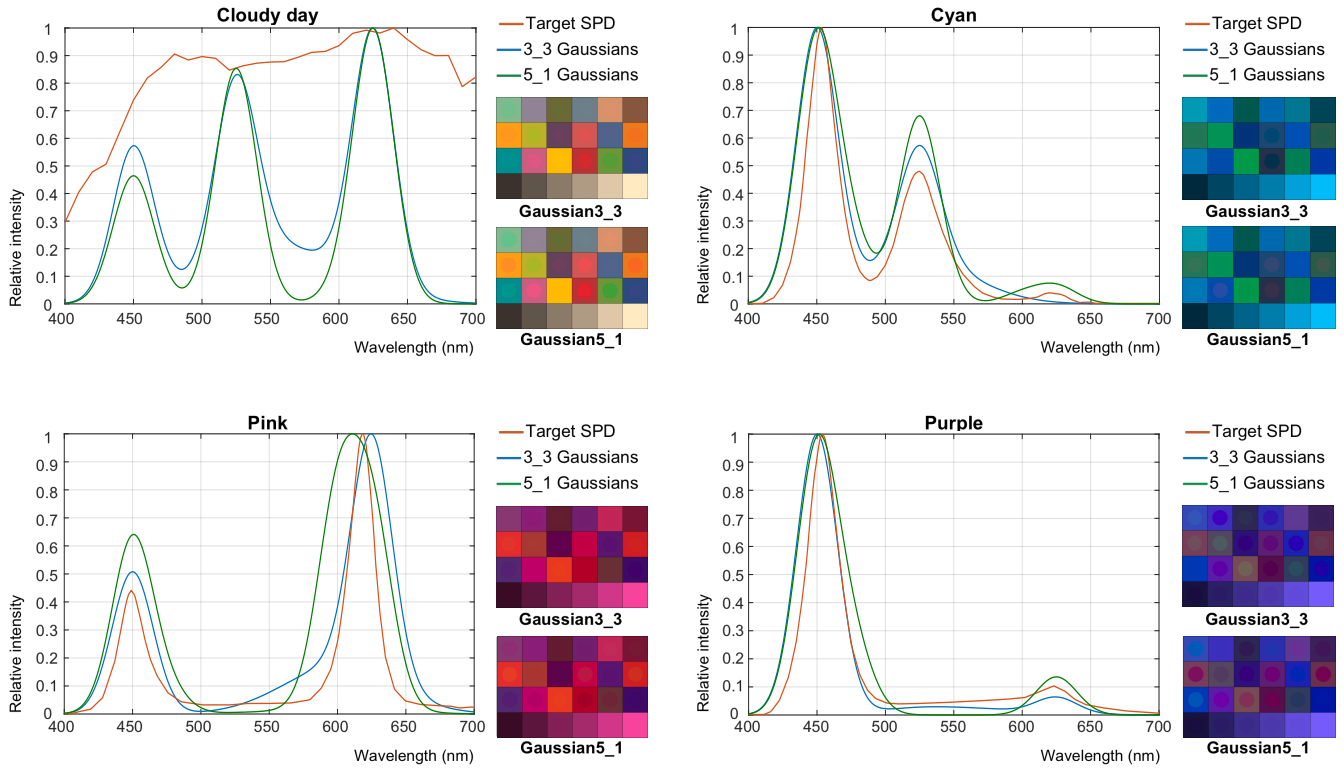


Figure 7: Additional spectral and color chart results for the selected optimal sets of Gaussians. Comparison shows the target data against the results of our method for 3 narrow and 3 broadband as well as the 5 narrow and 1 broadband Gaussian sets. We illuminate each color chart patch with the ground truth and the recovered spectra, and visualize the two colors in the same image: for each patch the inner circle corresponds to the color produced with our method and the outer frame represents the ground truth color

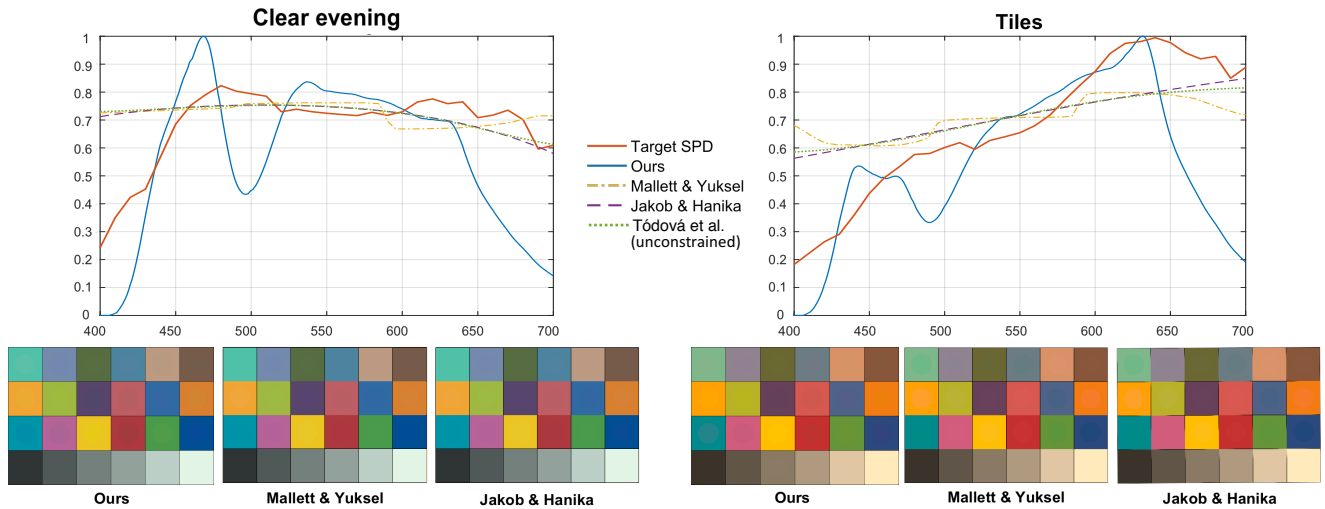
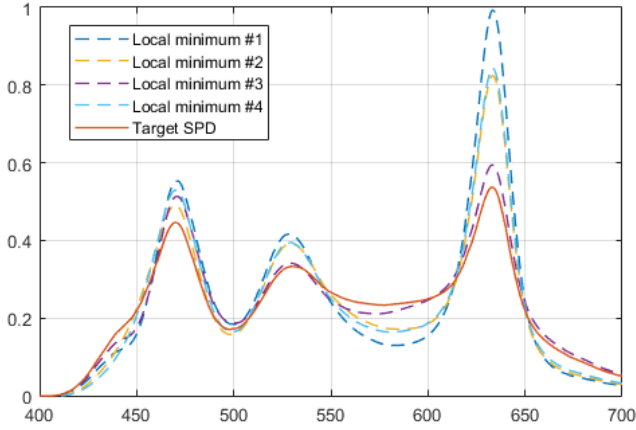


Figure 8: Additional comparisons to existing spectral upsampling methods for an outdoor illuminant (Clear evening) and a reflectance spectrum (Tiles). We use the tiles result as an illuminant spectrum to simulate its appearance in an environment map background pixel and show that for such use-case, our method recovers spectra with similar quality to previous work (on average, our method results in $\Delta E_{94} = 1.3657$ against $\Delta E_{94} = 1.7670$ for Mallett & Yuksel and $\Delta E_{94} = 1.7861$ Jakob & Hanika).



Spectrum	Weights	Relative XYZ values
Loc. min.#1	[0.7 0.8 0.5 0.1 0.9 0.0]	[0.630 0.635 0.510]
Loc. min.#2	[0.6 0.6 0.4 0.4 0.0 0.9]	[0.630 0.635 0.510]
Loc. min.#3	[0.3 0.5 0.5 0.8 0.9 0.0]	[0.630 0.635 0.510]
Loc. min.#4	[0.6 0.7 0.5 0.6 0.0 0.6]	[0.630 0.635 0.510]
Target SPD	[0.3 0.4 0.4 0.6 0.7 0.6]	[0.631 0.635 0.510]

Figure 9: Target spectrum (solid line) along with 4 local minima; such local minima lead to the same XYZ tristimulus values and hence same ΔE_{94} (0.02) distance from the target spectrum. Basis weights and XYZ values for all the spectra in this figure are reported in Table 2.

Table 2: Basis weights (1 decimal places) and XYZ values (3 decimal places) of the spectra reported in Fig. 9

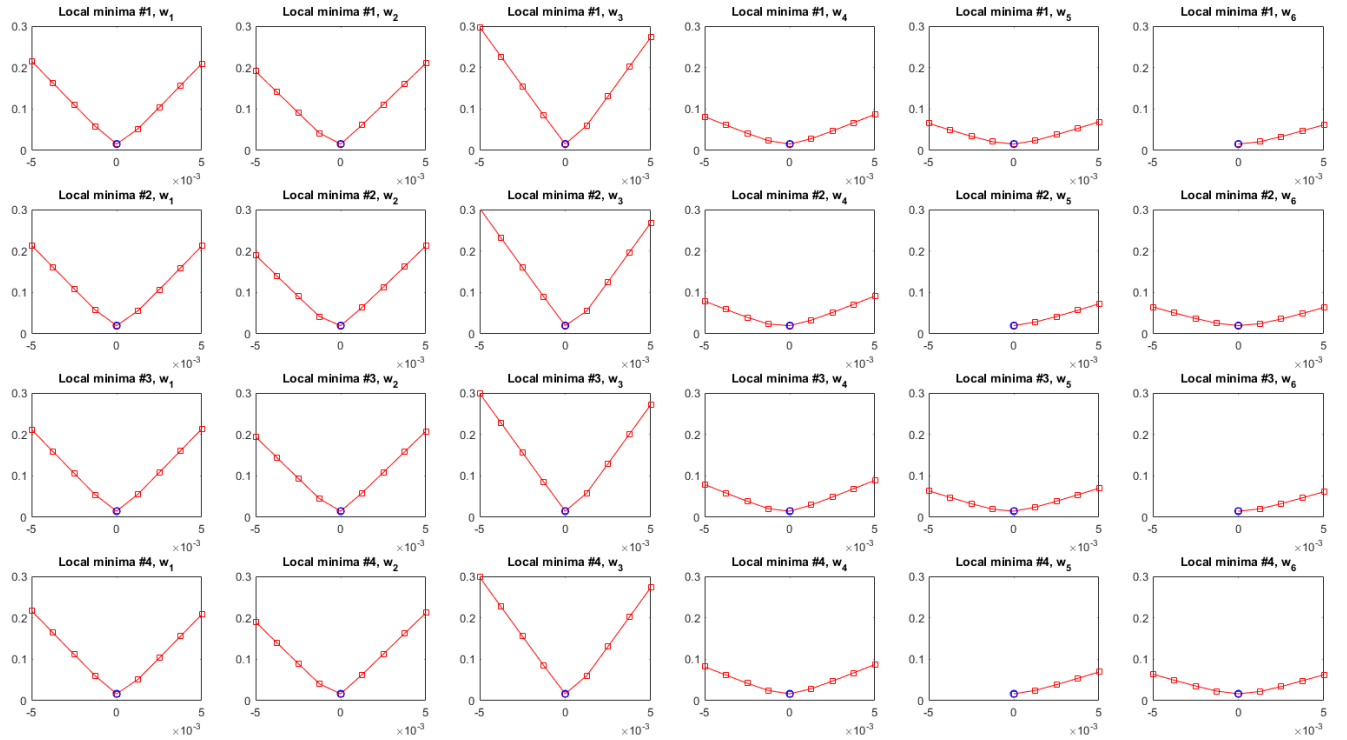


Figure 10: This figure demonstrates that the weights reported in Table 2 correspond to local minima. Rows refer to the local minima #1-#4 reported in Fig. 9 and Table 2, whereas each column refers to one of the weights. For each subplot, we allow one weight to vary in its local neighborhood ($[-0.005, +0.005]$, for values in the range $[0 - 1]$), while keeping the other 5 weights fixed to the values reported in Table 2. The horizontal axis reports the offset with respect to the weights indicated in the table, while the vertical axis shows the ΔE_{94} between the resulting spectrum and the target spectrum (reference). The blue circles highlight the weights reported in the table, while the red squares are the sampled neighbours. For all the examples, a change in any direction (by increasing or decreasing one of the weights) results in higher ΔE_{94} . The first 3 weights respectively corresponds to the narrowband LEDs red, green and blue, which lead to steeper changes in the objective function than the broadband ones (last three columns).

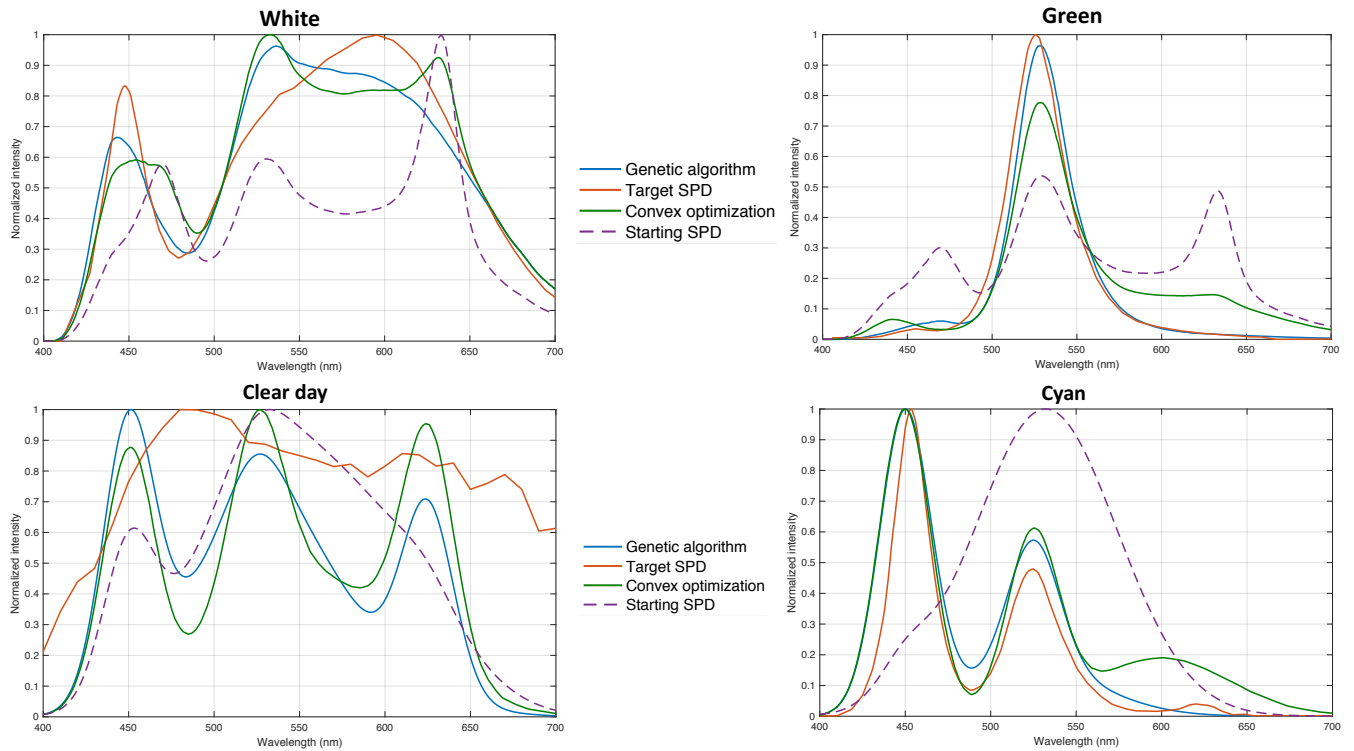


Figure 11: Comparison between the results achieved by the Genetic Algorithm optimization and by the Interior Point convex non-linear optimization (Convex optimization in the legend), both for our RGBW+ basis (top row) and for the Gaussian basis (bottom row). While convex optimization tends to produce optimized spectra that still preserve some features of the spectrum used for the initialization (barycentric interpolation, see Sec 3.2 in the main paper), on the average the stochastic nature of the GA allows to produce better solutions, less affected by local minima (i.e. with lower colorimetric error, often better matching the ground truth spectra, despite not being provided in input to the algorithms). In all cases, both algorithms have been compared on the same input, using the same initialization.




Figure 12: Comparison of full spectral rendering (middle row) of a scene lit with the cyan illuminant reported in Fig. 7, with rendering using the corresponding RGB values for the illuminant (top row), and rendering with spectral upsampling of RGB using our optimized Gaussian basis (bottom row). The result of our spectral upsampling is indistinguishable from the reference rendering, while color differences in the RGB rendering are clearly visible.