Supplemental Material: DIY hyperspectral imaging via polarization-induced spectral filters

Katherine Salesin, Dario Seyb, Sarah Friday, and Wojciech Jarosz

1 Introduction

This supplemental material provides further comprehensive analysis of our method. Given that our do-it-yourself method may be more prone to user error than setups with laboratory equipment in a controlled optical environment, we explore the sensitivity of our retrievals to user error in each stage of the measurement process in Sec. 2. We provide further context on how our method performs under different lighting scenarios in Sec. 3.

2 Parameter Sensitivity Analysis

Our method consists of three steps that could be contaminated by error:

1) Finding the birefringence of the tapes used to create the waveplates,
2) Taking real-world measurements of objects whose spectra we wish to retrieve, and
3) Reconstructing reflectance spectra from measurements.

Steps 1 and 2 contain possible sources of user or other real-world error that could infiltrate the results. The birefringence calibration step assumes that the thicknesses of the tapes are accurate, and any error introduced into the birefringence estimate by this one-time calibration will seep into reconstructions thereafter. Taking measurements, both for the birefringence calibration and for spectral reconstruction, requires rotating a stack of filters by hand to specific angles relative to a reference position, which introduces another possible source of error. Additionally, the camera's sensor will always inject some small amount of noise due to signal processing steps (read noise, quantization error), random fluctuation of photons across the sensor (photon shot noise), etc. Therefore, we chose to analyze the sensitivity of our method to errors in tape birefringence, tape thickness, angles of the analyzer and two waveplates in a filter stack, and sensor noise. In Figs. 1 to 4, birefringence is Δn, the angle of the analyzer is θ, and the angles of the two waveplates are α_1 and α_2, respectively.

• *K. Salesin, D. Seyb, S. Friday, and W. Jarosz are with Dartmouth College, Hanover, NH, USA.*
E-mail: katherine.a.salesin.gr@dartmouth.edu

2.1 Method

We ran virtual experiments simulating retrieval of the reflectance spectra of the Calibrite ColorChecker squares. In our virtual laboratory, we simulate taking measurements by multiplying the reflectance spectra of the ColorChecker squares by the transmission spectra of each of our filters. We then project these spectral products to RGB using the measured sensor responses of our Nikon D5100. Finally, we feed these virtual pixel measurements and the known transmission spectra of our filters as input to our reconstruction algorithm and retrieve the "unknown" reflectance spectra.

We simulated user error by adding zero-mean Gaussian white noise of increasing standard deviation to the six parameters individually (birefringence, tape thickness, angles of the analyzer and two waveplates, and sensor response). We also simulated user error in all parameters at once in the same manner, reflecting the real-world likeliness of some small amount of error permeating every parameter. We calculated the mean squared error (MSE) of each reconstruction by the formula $\mathrm{MSE} = \frac{1}{s}\sum_{i=1}^{s}\sum_{j=1}^{\lambda}(u_{i,j} - \hat{u}_{i,j})^2$, where s is the number of ColorChecker squares, u_i is the recovered value, and $\hat{u}_i$ is ground truth. We ran 500 trials of each noise-parameter set with different random seeds in order to analyze the spread of error.

We chose standard deviations of noise appropriate to each parameter being tested. For example, [1] state that even a 5 % difference in birefringence value can drastically effect the calculation of the phase difference caused by the waveplate. Therefore, in testing the birefringence, we chose standard deviations from 0–5 % of the birefringence. To test the tape thickness, we chose standard deviations of 0–2000 nm based on the reported [1] uncertainty of user-measured tape thickness. For all angle parameters, we chose standard deviations of 0–5.5° since the tick marks on the angle labels we used are in 5° increments. We estimated a realistic value for sensor noise by plugging the $\mathrm{SNR}_{\mathrm{dB}}$ 18% at ISO 100 for our camera [2] into the equation $\mathrm{SNR}_{\mathrm{dB}} = 20\log_{10}(\text{mean luminosity} \ / \ \text{std. dev. luminosity})$ and solving for the standard deviation of luminosity, then scaling from ADUs down to the range $[0, 1]$. We ran our simulations for both reasonable and extremely high noise levels to observe the effects of a range of RGB perturbations on our reconstructions.

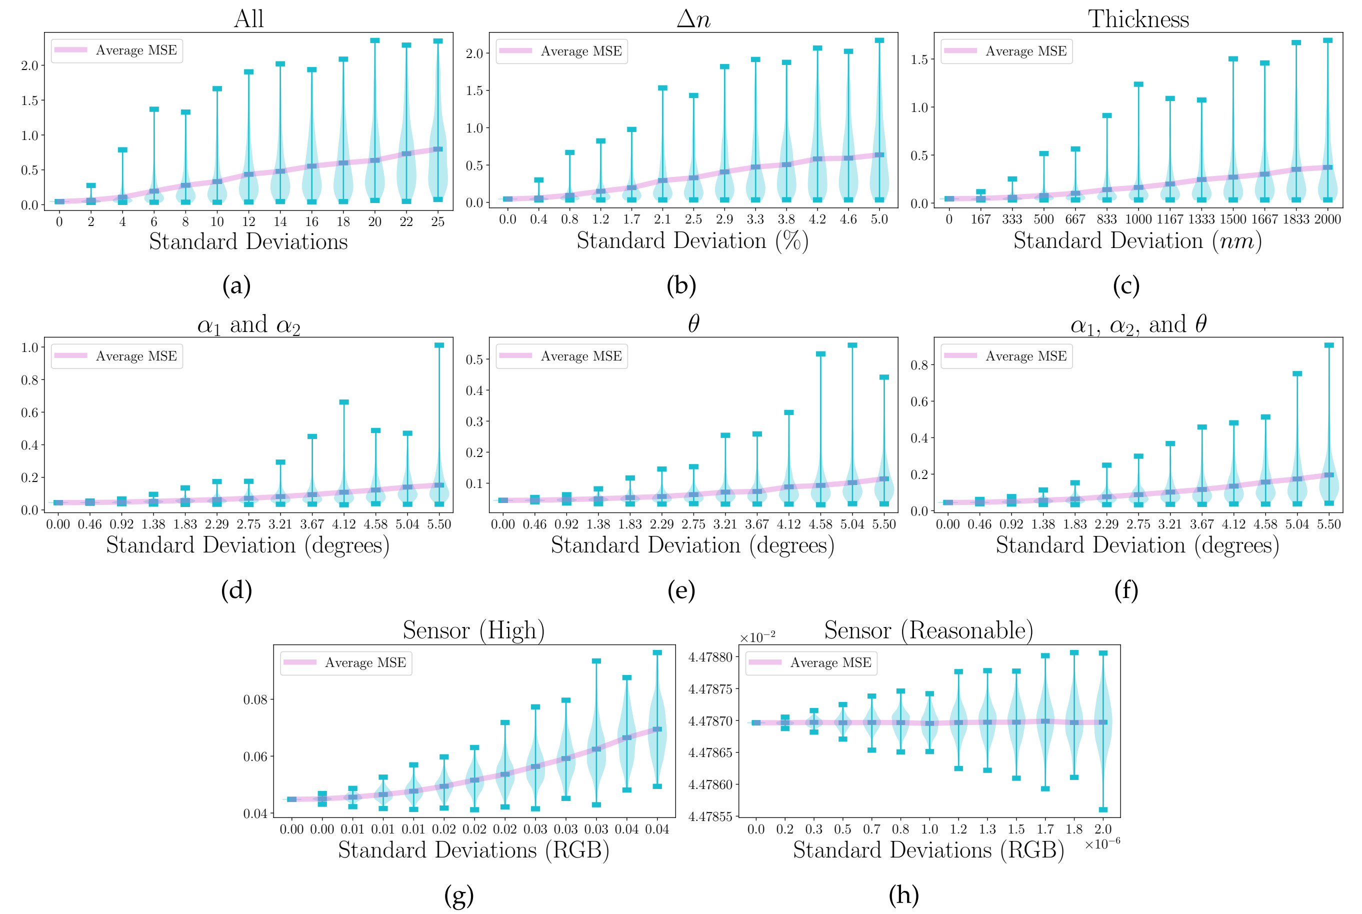

(a) (b) (c) (d) (e) (f) (g) (h)

Fig. 1: The distribution and average MSE of the reconstructed spectra increase gradually as higher levels of noise are added to birefringence (Δn), tape thickness, and filter stack alignment angles (α_1, α_2, and θ). (a) noise being introduced to all parameters at once. (b) 0–5 % noise being added to the birefringence. (c) 0–2000 nm of noise added to the tape thickness. (d) 0–5.5° of noise added to the waveplate angles. (e) 0–5.5° of noise added to the analyzer angle. (f) 0–5.5° of noise added to all filter stack angles. (g) 0–0.04 standard deviations of sensor noise applied to RGB sensor readings in [0,1], simulating unrealistically high noise. (h) 0–2.0×10^{-6} standard deviations of sensor noise, simulating more plausible noise.

2.2 Results

Fig. 1 shows how spectra reconstruction accuracy is affected by errors in the birefringence (Fig. 1b), thickness (Fig. 1c), and angles of the filters (Figs. 1d to 1f) across 500 trials. The violin plots show the distribution of the reconstruction error as kernel density estimations, while the line plot traces the mean MSE across noise intervals. As the noise increases for each parameter, the average MSE increases gradually without plateauing or spiking, which suggests that there is no particular threshold where the behavior of our system changes drastically.

Spectral reconstructions seem to be relatively insensitive to errors in filter angles (Figs. 1d to 1f) – even approaching the upper bound of 5.5°, which is unlikely in a real-world scenario since a user can reasonably estimate the angle to at least the nearest tick mark (within 2.5°). In practice, this means that being a degree or two off in alignment of the filter angles will not dramatically impact the spectral retrieval.

The reconstructions are also relatively insensitive to sensor noise (Figs. 1g to 1h), producing much less reconstruction error than other potential inaccuracies we tested. Therefore, one can expect that taking photos with any reasonable, modern camera under bright lighting conditions and a low ISO would introduce a negligible amount of reconstruction error.

In contrast, our method appears to be especially sensitive to errors in the birefringence (Fig. 1b) and thickness values (Fig. 1c) of the tape – adding noise to the birefringence results in the highest average MSE of any of the parameters alone. This suggests that minimizing error in the birefringence calibration step is especially important, as any inaccuracies will further plague the spectral reconstructions of later experiments relying on those estimates.

Figs. 2 to 4 visualize the reconstructed spectra at different noise levels, compared to ground truth and best-possible reconstructions. The color of each line is an sRGB projection of the recovered spectrum.

The best-possible baseline shown in Fig. 2, in which the parameters have zero noise, illustrates the lowest error that our reconstruction method can achieve. The resulting spectra closely match the ground truth spectra. However, there is still some error, especially at the outer limits of the wavelength range. The error pattern closely matches that observed in our real-world ColorChecker reconstructions

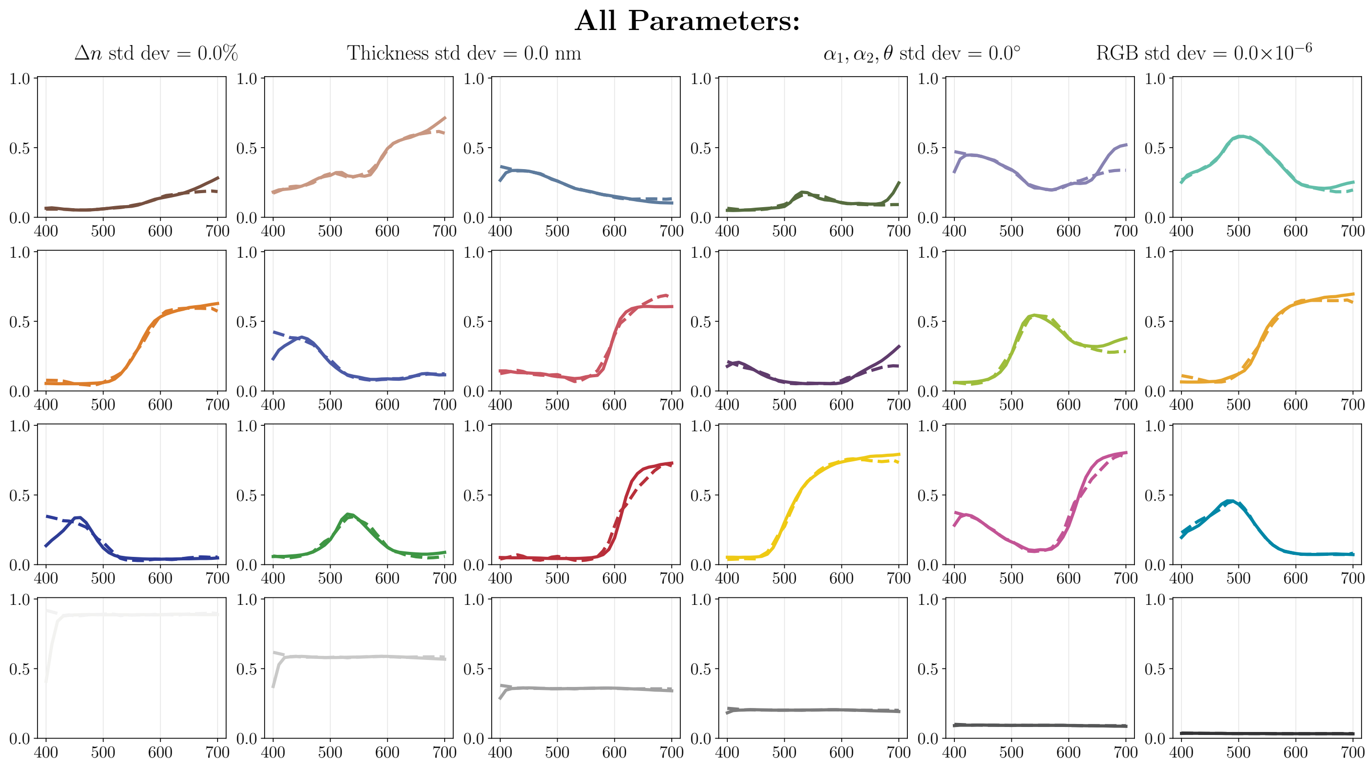

Fig. 2: The reconstructed ColorChecker spectra (dashed) under ideal conditions (no errors in the birefringence Δn, thickness, or angles) shows good agreement to the ground truth (solid). The colors of each line are the sRGB projections of the spectra

(Fig. 6 in the main paper) and in prior linear-reconstruction-based methods [3], suggesting that this error is inherent to the problem formulation itself – due to the limitations of representing the various sets of spectra and spectral responses by sets of basis functions, and also the available information for reconstruction tailing off to zero at the outer limits of the sensor responses.

In Figs. 3 to 4, as noise is added to the input parameters, the error in reconstruction clearly manifests as flaring and pinching at particular wavelengths. We postulate that this flaring and pinching is due to the representation of the reflectance spectra by a particular set of basis functions. Notably, our method will generally retrieve a metameric spectrum that maps to the same RGB color as the ground truth spectrum, although it struggles with light grays in particular.

3 Lighting Conditions & White Balance

In Fig. 5, we examine the performance of our reconstruction method on retrieving the ColorChecker spectra under three real-world lighting environments: a cloudy day, a sunny day, and a room with warm LEDs.

In each setting, the white-balanced retrievals far outperform the un-white-balanced retrievals in terms of accuracy. In comparing the white-balanced retrievals to one another, they perform similarly (SSE: 2.67 / 2.81 / 3.84, MRSE: 0.148 / 0.119 / 0.171 for sunny daylight/cloudy daylight/LEDs), generally showing the same error pattern as described above in Sec. 2 with some minor variations.

The estimated sensor responses are generally consistent across lighting conditions. The estimated illuminants, however, clearly deviate from their ground truth emission spectra. We postulate that there is not enough information available from the known ColorChecker reflectance spectra to disambiguate the signal between the illuminants and sensor responses – and so the sensor responses and illuminant spectra end up trading off whatever part of the signal happens to be easiest to represent in their respective sets of basis functions. For example, in Fig. 5 it is clear for the LED setting that some of the LED emission peak around 450 nm is being baked into the green sensor response. This should not matter in theory for spectral reflectance retrieval since it is only the product of the sensor responses and the illuminant that matters for reconstruction, but as described in Sec. 6.3 of the main paper, it could be an issue for some applications where very accurate estimation of the sensor responses or illuminant spectrum *in addition to* the reflectance spectra is required. In practice, we found that increasing the number of basis functions used to represent the sensor responses and illuminant by a few – beyond the number that could encompass the vast majority of the variance of their respective datasets – helped to inject a bit more flexibility in the system and decreased the error of the reconstructions.

References

[1] A. Beléndez, E. Fernández, J. Francés, and C. Neipp, “Birefringence of cellotape: Jones representation and experimental analysis,” *European Journal of Physics*, vol. 31, no. 3, mar 2010.

[2] DxOMark. Nikon d5100. [Online]. Available: https://www.dxomark.com/Cameras/Nikon/D5100---Measurements

[3] S. W. Oh, M. S. Brown, M. Pollefeys, and S. J. Kim, “Do it yourself hyperspectral imaging with everyday digital cameras,” in *Proc. IEEE CVPR*. New York, NY, USA: IEEE, 2016.

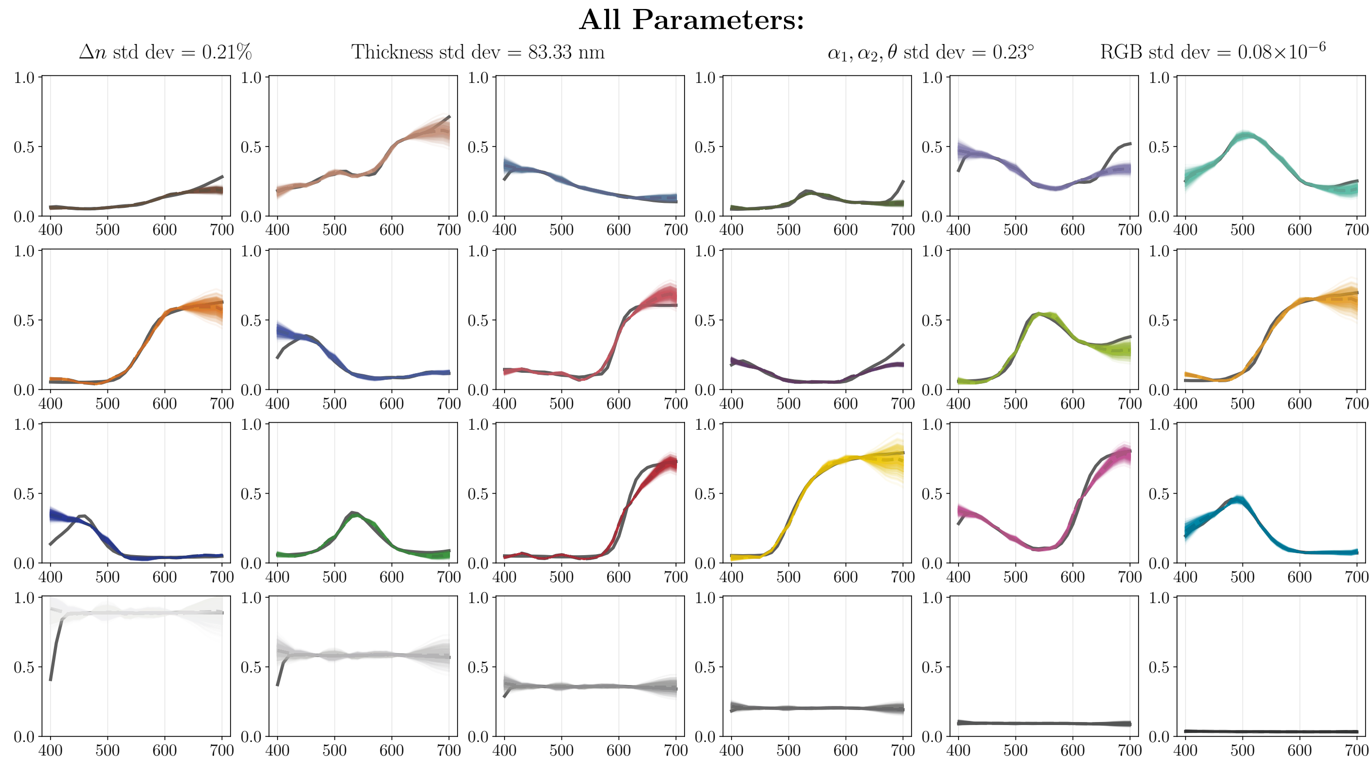

Fig. 3: The reconstructed ColorChecker spectra flare out from the ground truth (solid gray) and ideal reconstruction (dashed gray) at the ends of the sensor response wavelengths as one interval of noise is introduced to all parameters at once.

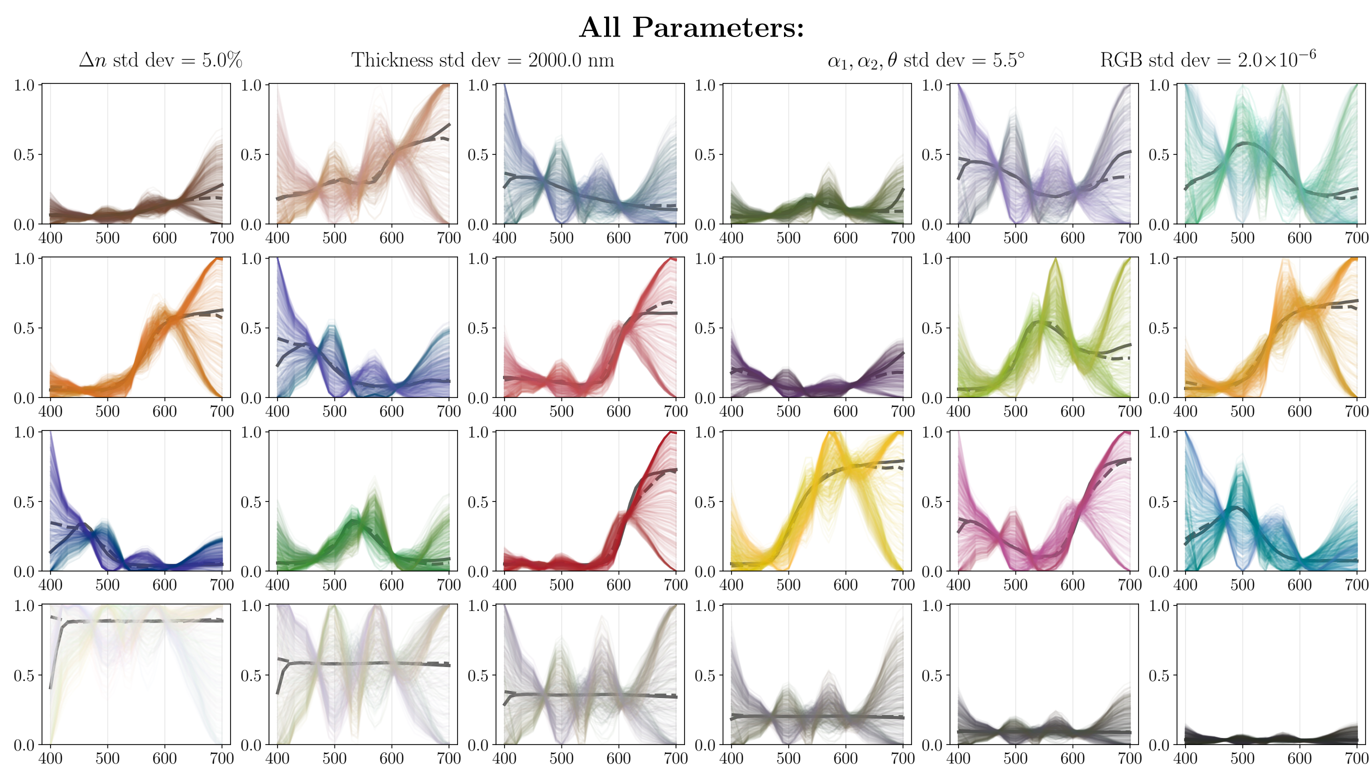

Fig. 4: Adding the maximum amount of noise results in wild flaring and distinct pinching at specific wavelength intervals. In general, the noisy reconstructed spectra map to similar sRGB values as the ground truth (solid gray) and ideal reconstructed (dashed gray).

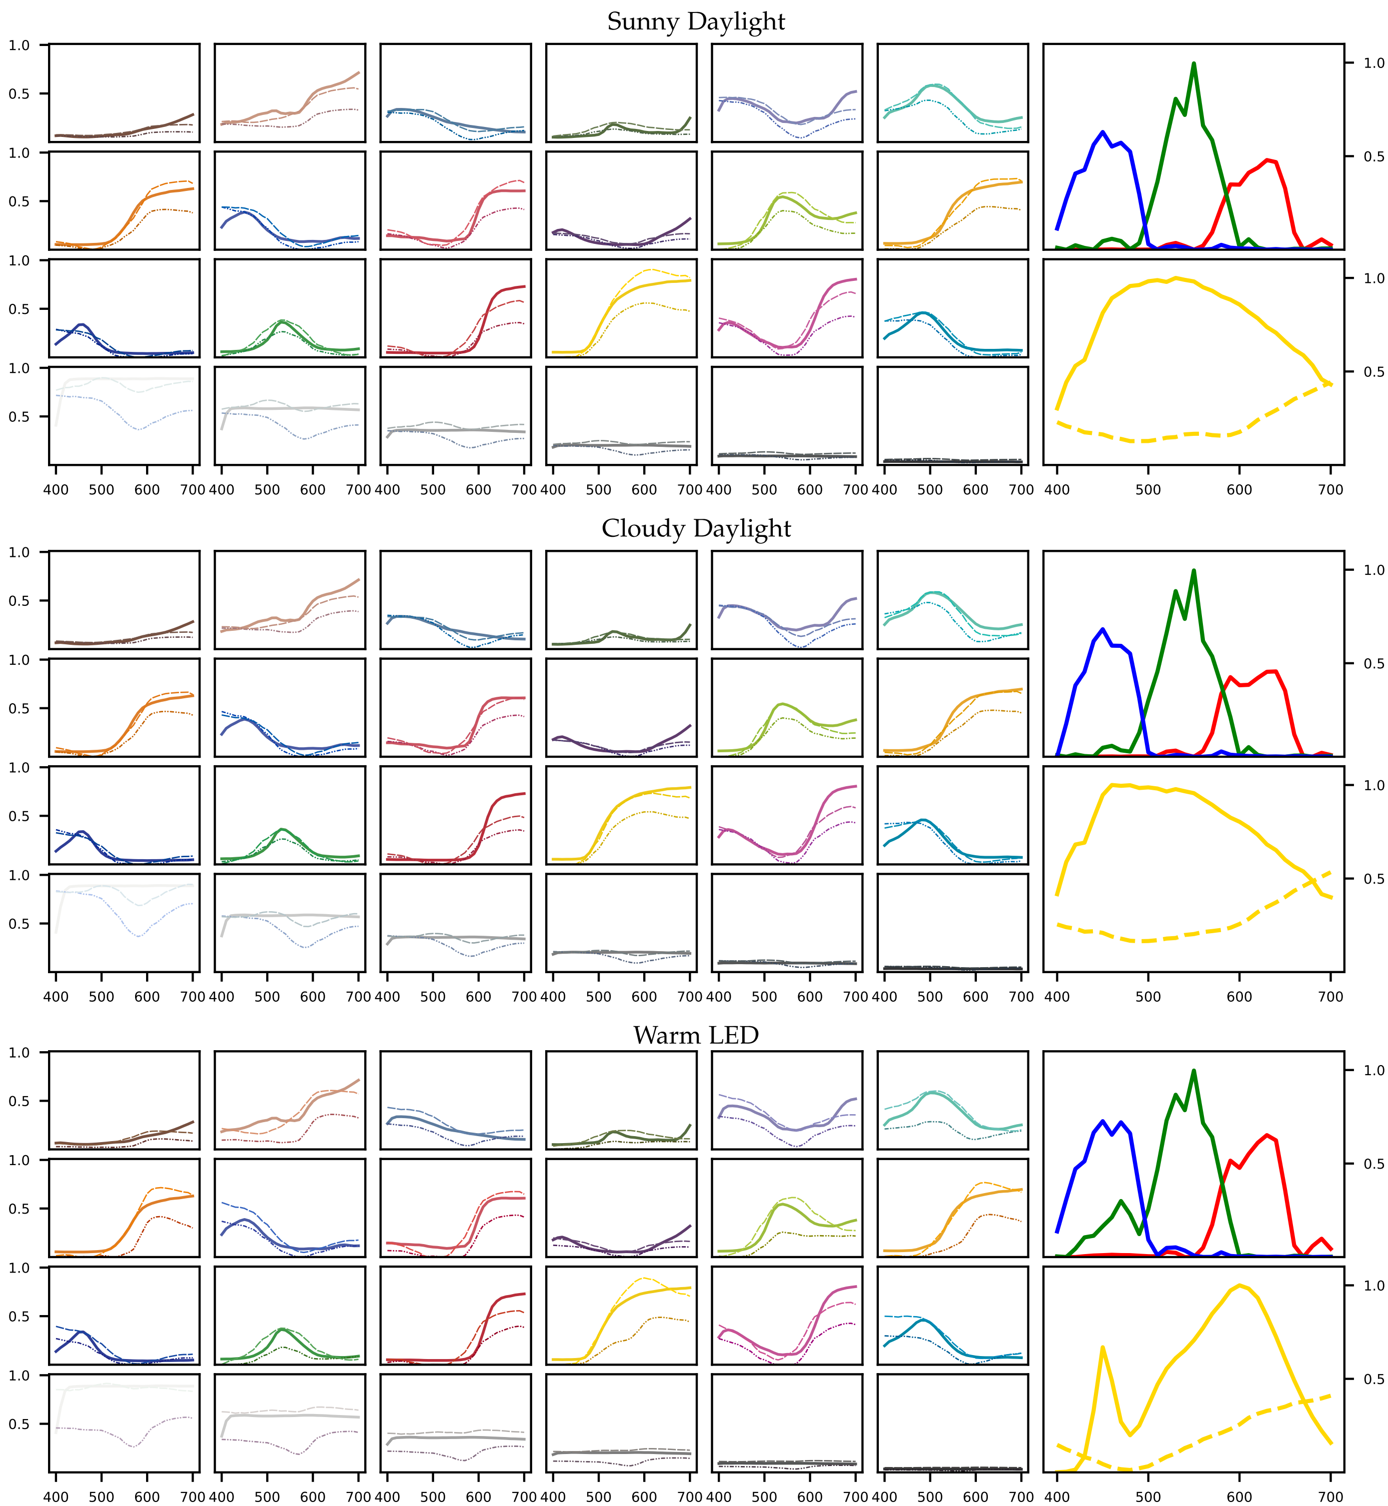

Fig. 5: Additional real-world reconstruction results of our method in various lighting conditions. In addition to the retrieved ColorChecker reflectance spectra, for each setting we show the estimated sensor responses and illuminant spectra (solid: ground truth, dashed: estimated) retrieved by our white balance step.