

Groundtruth Data for Multispectral Bidirectional Texture Functions

Martin Rump, Ralf Sarlette and Reinhard Klein
Universität Bonn, Germany

Abstract

In this paper we present a method for the acquisition of multispectral Bidirectional Texture Functions (BTFs) for the color-correct reproduction of complex materials in virtual scenes with arbitrary illumination. We explain our acquisition setup in detail as well as its careful calibration. The main purpose of the described setup is to provide accurate groundtruth data to verify future methods for simplified multispectral BTF measurements as well as techniques for efficient compression and rendering of multispectral BTF data. We show the accuracy of data obtained with our setup by comparing rendered images to photographs of real scenes. We also verify our datasets using reference measurements obtained with a spectrometer.

Introduction

Predictive rendering is of great importance in many parts of the industry for the use in virtual prototyping or virtual showrooms. Here, enormous progress has been made in the visual reproduction of real materials by measuring optical material properties. But until now, in computer graphics most effort is spent on RGB or similar tristimulus representations. Due to metamerism, tristimulus based reflectance data cannot be used to reproduce the measured material under arbitrary lighting conditions and is thus not sufficient for predictive rendering.

Therefore, fully spectral measurements are necessary. Multispectral photography has been used to reproduce colors for a long time but devices for the acquisition of bi-directional reflectance data are so far limited to homogeneous materials represented by the Bidirectional Reflectance Distribution Function (BRDF). However, for many materials these representations are not sufficient for an accurate reproduction under arbitrary illumination conditions. Two examples for such complex materials can be seen in Figure 1. The handmade color checker does not only contain mainly diffuse color fields, but also two very specular fields made of silver and gold paint that furthermore contain sparkling particles. The red fabric consists of four different kinds of basic strings (silver, green and two kinds of red) and has a very complex mesostructure. Due to this structure it shows very strong anisotropic effects. For these materials, homogeneous reflectance data or single viewpoint photographs will not be sufficient for an accurate reproduction.

The Bidirectional Texture Function (BTF) was first introduced by Dana et al. [3] and can be understood as a series of pictures of a surface under varying illumination and viewing angles. Since these pictures already contain effects like self-shadowing, occlusion, interreflection and subsurface scattering, the BTF can represent a broad class of materials with spatial variations in reflectance and with nearly arbitrary complex mesostructure. BTFs are thus able to represent materials, where BRDFs or even more simple measurements from single viewpoints are not sufficient.

The BTF is a six-dimensional function dependent on illumination and viewing direction ω_i, ω_o as well as the surface point x :



(a) Handmade color checker. Fields row-wise: White, Paper, Fluorescent Yellow, Fluorescent Red, Turquoise, Blue1, Blue2, Blue3, Gold, Silver, Green1, Green2, Yellow, Orange, Red1, Red2



(b) Red fabric with strong anisotropic effects. Rotated 90° from left to right image.

Figure 1. The two test materials captured with the setup

$\rho_{BTF}(x, \omega_i, \omega_o)$. The bi-directional reflectance of a single surface point of the BTF is called *Apparent-BRDF* since it is very similar to a BRDF but already contains the various effects described above.

While efficient devices for the capture of RGB-BTFs are already at hand, multispectral measurements of BTFs have not been made until now. This is due to several reasons. First, specialized and costly hardware is required, whereas RGB measurements require only simple RGB cameras. Second, the use of narrow band-pass filters severely reduces the light arriving at the sensor, leading to a bad signal-to-noise ratio. In contrast to BRDF measurements, the noise cannot be filtered out by averaging over many sensor pixels, since spatially varying reflectance information is to be captured and thus every single pixel must be treated individually. Third, the amount of captured data is critical. If the spectral sampling is high enough to reproduce high-

frequent features in the reflectance spectra, the data will easily become ten times larger than RGB data with the same angular and spatial resolution.

In this paper, we present an approach for the measurement of multispectral BTF. The focus of this work was to create highly accurate data accepting high measurement times. For now, we restricted ourselves to materials without fluorescence and phosphorescence effects because the additional wavelength dimension would lead to unacceptable amounts of data and extremely long measurement times. Furthermore, the amount of light arriving at the sensor would be reduced by a further filter and thus the exposure times must be increased to compensate for this leading to even higher measurement times.

Due to the long measurement times, the setup is not suitable for the acquisition of whole material databases but it provides groundtruth data that can be used for future research on fast measurement, compression and efficient rendering of multispectral BTFs. Therefore, after publication we will make the data available at the BTF Database Bonn [1]. We demonstrate the accuracy of the data by comparing with spectrometer measurements and by comparing rendered images with multispectral photographs.

Previous Work

Spectral BRDF measurements Multispectral BRDF measurements have been made for a long time and measurement devices are already at hand. These are either goniometer based devices (e.g. [9], [10]) or devices with multiple sensors (e.g. [2]). A device like the one from Marschner et al. [6] that captures curved objects and transfers spatial resolution to angular resolution could also be equipped with a multispectral camera.

The device from Tsuchida et al. [10] should be capable of capturing a multispectral BTF, but they showed only the application to BRDF. Furthermore, the measurement speed of their setup is quite low due to the mechanical filter wheel in front of the light source.

RGB BTF measurements A first measurement device for RGB-BTF was already presented by Dana et al. [3]. They took 205 images per material using a RGB video camera. The devices by Sattler et al. [8], Koudelka et al. [5] and McAllister et al. [1] are very similar, but due to full automation they are able to capture a much larger amount of view- and light directions. Due to the low cost of RGB consumer cameras, also fully or at least partially parallelized measurement setups are possible. A description of such a device can be found in [7]. Unlike for BRDFs, multispectral measurements of BTFs have not been made until now.

Measurement Setup

We basically extended the setup of Sattler et al. [8] with a monochrome camera and a liquid crystal tunable filter (LCTF) for multispectral measurements (see Figure 2 for a photograph).

Since only low amounts of light pass the bandfilter, we needed a camera with a very good signal-to-noise ratio. We therefore use a Photometrics CoolSnap 4k because it is cooled to approx. -25°C effectively reducing the noise level. It has 2048×2048 pixels resolution, a 12bit ADC and has a highly linear response curve (see next section for details).

The bandpass filter is a CRI VariSpec VS10, which can be tuned to wavelengths between 400 and 720 nm. The spectral transmission curve of the filter is gauss-shaped with 10nm FWHM. Due to the absence of mechanical parts it allows for

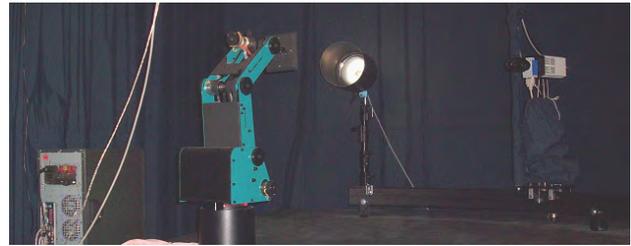


Figure 2. The measurement setup

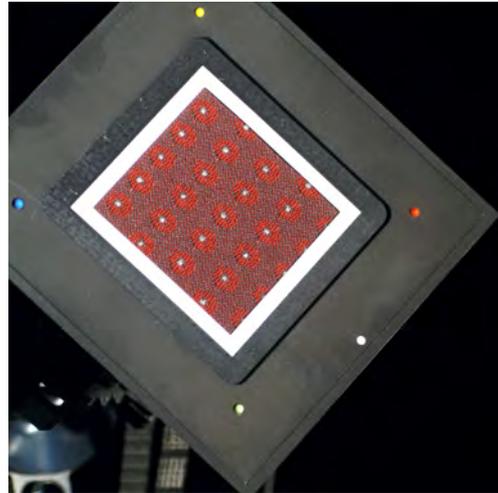


Figure 3. One photograph of the red fabric BTF.

very fast tuning in the order of 50ms. This is in stark contrast to the use of a mechanical filter wheel ([10]).

Since the bandfilter is based on interference effects, it must be ensured that the incident light is nearly parallel. The filter has an acceptance angle of 7° to the optical axis. To ensure, that light rays reflected from the target into the camera do not have a greater incident angle, an optical system from Schneider Kreuznach is used that ensures that these light rays enter the filter nearly parallel to the optical axis. Despite the use of this optical system, at the corners of the image the rays are not fully parallel to the optical axis leading to a slight darkening in these regions. This can be neglected since we only extract the target area from the photographs, which is centered in the images (see Figure 3).

The lamp was a 570W HMI lamp with a UV filter to prevent damage of the target material.

Calibration

In this section we explain all calibration steps for the measurement setup. They concern the camera and filter as well as the lamp.

First of all, the spectrum of the light source was measured using a X-rite I1 handheld spectrometer. The result can be seen in Figure 4.

To measure radiance values with the camera, *dark frames* must be subtracted from the images to correct for hot pixels of the camera. For every wavelength band a dark frame is captured using the same exposure times as in the BTF measurement. Second, the response of the camera must be inverted to get energy values from the cameras pixel values. For this, we determined an inverse response curve of the camera for every wavelength band by linear interpolation between 200 measurement points generated by taking shots of a white standard with varying exposure time. The result of the inverse response after darkframe subtrac-

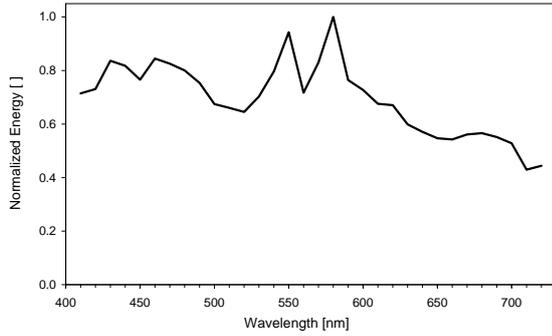


Figure 4. Spectrum of the HMI lamp

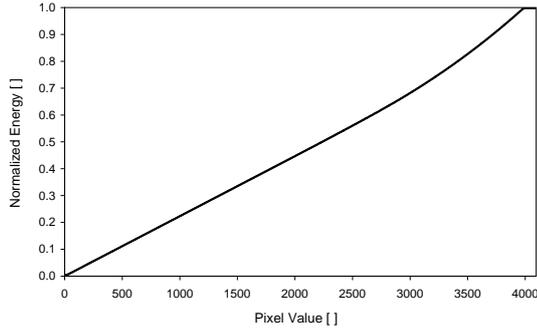


Figure 5. Inverse response of the camera

tion is shown in Figure 5. It can be seen that the response of the camera is mostly linear and shows only little saturation effect at higher pixel values. An energy value can be derived from a pixel value $P_{x,\lambda}$ by applying the following formula:

$$E_{x,\lambda} = RC_{\lambda}^{-1}(P_{x,\lambda} - D_{x,\lambda}) \quad (1)$$

with the response curve RC_{λ} and the corresponding dark-frame value $D_{x,\lambda}$.

Since the LCTF has a different opacity for each wavelength band and the cameras pixels have a different sensitivity as well, we had to determine a scale factor per wavelength band to correct for this. Furthermore, the conversion from energy values to radiance values requires an additional constant factor. We measured them at once by taking photographs at every wavelength of a white standard made of barium sulfate. It has a known diffuse reflection spectrum W_{λ} with approx. 97% albedo. The white standard was placed at the measurement target, with the light direction perpendicular to the white standard. The camera and light source were at the same distance as we used for the measurement. The calibration factors were calculated as the ratio of the linearized response of the camera/filter E_{λ} to the known spectrum of the light source I_{λ} reflected by the known spectrum of the white field W_{λ} integrated over the exposure time T :

$$F_{\lambda} = \frac{I_{\lambda} W_{\lambda} T}{\pi E_{\lambda,white}} \quad (2)$$

where the energy $E_{\lambda,white}$ has been determined as a mean energy over the values $P_{x,\lambda,white}$ of the N pixels containing the white field in the photograph:

$$E_{\lambda,white} = \frac{1}{N} \sum_{x=1}^N RC_{\lambda}^{-1}(P_{x,\lambda,white} - D_{x,\lambda}) \quad (3)$$

The result from this calibration can be taken from Figure 6.

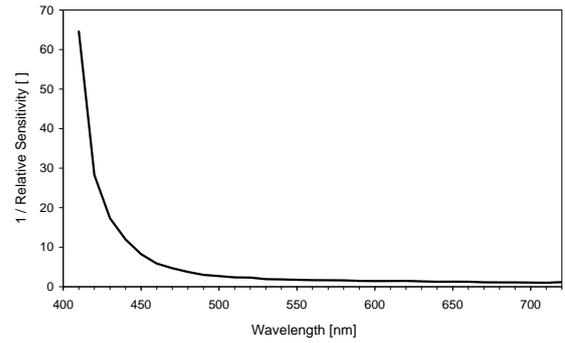


Figure 6. Sensitivity of the camera/filter combination in different wavelength bands

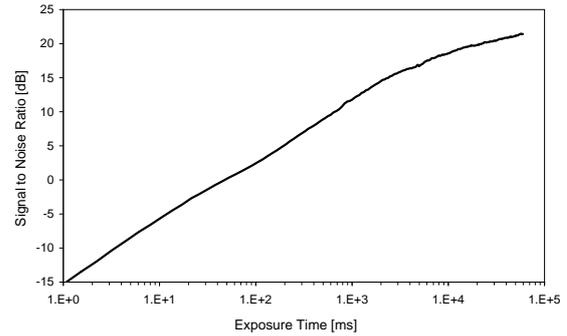


Figure 7. Signal-to-noise ratio of the camera for increasing exposure times. It has been measured with a very dark illumination (60 Lux halogen lighting) at 600nm using the 10nm wide LCTF.

As it can be seen, the combination of camera and filter has very low sensitivity in the bands of the visible spectrum that correspond to violet colors. This is a quite general problem, since the measured dynamic range of the captured scene is extended due to the properties of the capture device. There are two possible ways to compensate for this. First, one could simply scale the resulting energies by the measured factor. Second, the exposure times T can be increased corresponding to the falloff in sensitivity to be $T_{\lambda} = F_{\lambda} T$. To judge which approach is better, we measured the signal-to-noise ratio of our camera over many shutter times. The result can be seen in Figure 7. It shows, that the SNR of the camera steadily increases with the exposure time. This is true even for very long exposures of about 50 seconds. This enabled us to still capture stable data at low wavelengths by increasing the shuttertime as described above. If the energy values are scaled instead, the SNR ratio would be very bad and thus the data not reliable.

Having obtained the calibration factors, the final radiance value of one pixel x of the BTF is then calculated as follows:

$$L_{x,\lambda} = \frac{F_{\lambda} RC_{\lambda}^{-1}(P_{x,\lambda} - D_{x,\lambda})}{I_{\lambda} T_{\lambda}} \quad (4)$$

Measurement and Processing

Measurement process We decided to sample the visual spectrum from 410nm to 720nm in 10nm steps acquiring 32 monochrome images per light and view direction. The sensitivity of the camera/filter combination at 400nm is very low leading to long exposure times to get reliable results. Since this would nearly double the total measurement time and since these band is perceptually not very relevant we excluded it for now and want to integrate it again as soon as a stronger light source is available.

We used the same angular sampling as Sattler et al. [8] and thus capture $81 \times 81 = 6561$ view and light directions.

All parts of the setup are controlled by a single PC. The necessary steps to acquire each image of the BTF are:

1. Move robot and camera to correct position
2. Tune LCTF to desired wavelength
3. Take shot with camera
4. Save captured image
5. Continue at step 2 until all wavelength bands are sampled

Step 1 takes about 1-5 seconds if the robot arm is to be moved only and much longer if the camera must be moved on the rail to its next position. To shorten the measurement time, the captured images are ordered in a way, that the camera movements are minimized. Steps 2 - 5 take about 28 seconds for all 31 bands if the saving and shooting steps are parallelized. The total exposure times are 12 seconds, the download times are 12.5 seconds and the tuning of the filter takes about 2s. The 1.5 additional seconds are spent on camera control since the camera is run in single shot mode. Since we capture 81×81 multispectral images, we end up with a total measurement time of 60 hours (2.5 days). For now, we did not take multiple shots with different exposure times per direction pair because otherwise the measurement time would be too long. This will be possible in the future since we want to integrate a stronger lamp and can then reduce the exposure times. Furthermore, we want to parallelize the image download from the camera and the tuning of the LCTF and are then capable of using the continuous operation mode of the camera. We hope that it will allow us to measure an additional exposure step within approximately the same time.

Post-Processing Since the camera has a 12bit ADC, every single 2048×2048 image requires 6MB space if 2 pixels are packed in 3 bytes. We therefore end up with 81 lights \times 81 views \times 32 raw images \times 6 MB = 1.2 TB of raw data. The raw pixel values are converted into radiance values in a postprocessing step. Here, the inverse response curve and the calibration factors are applied as explained in equation 4 and the 32 images corresponding to one multispectral BTF image are assembled into one file. We used the OpenEXR format here, because it supports floating point data and an arbitrary number of channels. We end up with approx. 1.2 TB of EXR image files. Afterwards the white corners of the target are detected in all images and the target area is rectified to register the images in the texture coordinate system. Here, a resolution of 800×800 pixels is used which approximately corresponds to the target size in the frontal views. This corresponds to 485 GB of final data.

Compression Since the amounts of captured data are not suitable for the use in a rendering system, we had to apply a compression method. We used a SVD based compression method that is a modification of the method first described in [5]. We first split the spectra into a brightness channel and normalized spectra and factorized the two parts individually. The brightness channel contains most of the complexity that arises from the mesostructure of the BTF. For this reason, we used 250 PCA components to represent it, whereas 60 components were sufficient to reproduce each single band of the spectra. For a BTF with 740×740 texels and 81×81 directions and 31 bands we end up with approx. 2 GB of compressed data without a noticeable loss in accuracy.

Analysis of spectra SVD could also be applied to the normalized spectra before the compression deriving an eigenbasis for



Figure 8. Environment map of a room with complex lighting consisting of large neon lamps and small halogen spotlights.

the material consisting of characteristic spectra. If the material has only a small number M of different basic spectra, the singular values after the M -th one should drop near zero and thus a further dimensionality reduction would be possible.

Results

To show the accuracy of our datasets we made several comparisons. First, we compared the spectra of different parts of the material samples measured with our setup to corresponding measurements using the I1 handheld spectrometer. The I1 illuminates a circular spot of the sample with an angle of $\theta = 45^\circ$ from all sides and detects the reflected spectrum from the top. To simulate this, we integrated our data over circular spatial regions as well as all the images of the BTF that correspond to illumination directions of $\theta = 45^\circ$ and viewing directions from the top. Since we do not know the exact geometry of the spectrometer we cannot hope to reproduce all of the data exactly. This is especially true for the specular fields since here a small angular deviation results in strong differences in brightness. The comparison result can be seen in Figures 10 and 11. Our setup is able to reproduce the spectra with very high precision. The fluorescent colors on the colorchecker cannot be reproduced which results from the different amount of energy emitted by the light source of the I1 and our HMI lamp in the UV range. Since we did not sample the dimension of wavelength of incoming light, this effect cannot be reproduced.

For the second comparison we took a multispectral photograph of the color checker in a room with complex lighting consisting of small thermal light sources (halogen lamps) as well as large fluorescent light sources (neon lamps). The resulting illumination spectrum has sharp peaks and is quite different compared to typically used daylight illuminants.

Then we captured the illumination conditions extending the method from Debevec et al. [4] by taking multispectral photographs of a chrome sphere mounted at the position of the material samples. The environment map is shown in Figure 8.

Using this image based lighting data and the captured BTF data we re-rendered a photograph of the color checker using a spectral rendering system. A comparison of the rendering and the photograph can be seen in Figure 9. The multispectral images have been converted to RGB by convolving the spectra with the CIE 1931 RGB color matching functions whereas the environment map and the BTFs for the RGB comparison were converted to RGB before the rendering. All three images have been white balanced using the same RGB factor derived from the environment map and tonemapped using a measured response curve of a RGB camera afterwards. Calibrated tonemapping and display of the multispectral images is out of scope of this paper and thus not further discussed. The multispectral rendering reproduces the

photograph faithfully and shows the accuracy of our reflectance data. The third image of Figure 9 shows that RGB BTFs cannot be used for predictive rendering due to metamerism effects.

Conclusions

We presented our effort for the construction and calibration of a multispectral BTF measurement setup. Furthermore, we showed the accuracy of our groundtruth data by comparing it to spectrometer data and by comparing a rendered image with a photograph. We will make the materials available for download to enable other researchers to use the data for evaluation of rendering and compression methods. In future work we plan to measure more materials and increase the dynamic range of the data. For even better quality of the data we will integrate a more powerful QTH light source with a very smooth spectrum. Furthermore, the compression of the spectral BTF data should be analyzed in more detail. For the practical application of spectral BTF, more efficient measurement techniques must be elaborated, which may be checked against our groundtruth database.

Acknowledgment

This work was supported by the German Science Foundation (DFG) under research grant KL 1142/7-1.

References

- [1] Btf database bonn. <http://btf.cs.uni-bonn.de>.
- [2] Moshe Ben-Ezra, Jiaping Wang, Bennett Wilburn, Xiaoyang Li, and Le Ma. An led-only brdf measurement device. In *CVPR*. IEEE Computer Society, 2008.
- [3] Kristin J. Dana, Bram van Ginneken, Shree K. Nayar, and Jan J. Koenderink. Reflectance and texture of real-world surfaces. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 151–157, 1997.
- [4] Paul Debevec. Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *SIGGRAPH '98: Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 189–198, New York, NY, USA, 1998. ACM.
- [5] Melissa L. Koudelka and Sebastian Magda. Acquisition, compression, and synthesis of bidirectional texture functions. In *In Proc. 3rd Int. Workshop on Texture Analysis and Synthesis (Oct)*, pages 59–64, 2003.
- [6] Stephen Marschner, Stephen Westin, Eric Lafortune, Kenneth Torrance, and Donald Greenberg. Image-based BRDF measurement including human skin. In *Eurographics Workshop on Rendering*, pages 131–144, 1999.
- [7] M. Rump, G. Müller, R. Sarlette, D. Koch, and R. Klein.

Photo-realistic rendering of metallic car paint from image-based measurements. *Computer Graphics Forum*, 27(2):527–536, August 2008.

- [8] M. Sattler, R. Sarlette, and R. Klein. Efficient and realistic visualization of cloth. *Proceedings of the Eurographics Symposium on Rendering 2003*, 2003.
- [9] G. Serrot, M. Bodilis, X. Briottet, and H. Cosnefroy. Presentation of a new BRDF measurement device. In A. D. Devir, A. Kohle, U. Schreiber, & C. Werner, editor, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 3494 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 34–40, December 1998.
- [10] Masaru Tsuchida, Hiroyuki Arai, Masami Nishiko, Yoshiyuki Sakaguchi, T Uchiyama, Masahiro Yamaguchi, Hideaki Haneishi, and Nagaaki Ohyama. Development of brdf and btf measurement and computer-aided design systems based on multispectral imaging. In *Proceedings of AIC05*, pages 129–132, 2005.

Author Biography

Martin Rump studied computer science at the University of Bonn, Germany, from where he received his diploma in computer science (Dipl.-Inform.) in 2008. Since then he has worked in the Computer Graphics Group at University of Bonn, Germany. His work has focused on realistic materials in computer graphics.

Ralf Sarlette studied physics at the RWTH Aachen University, Germany, from where he received his diploma in physics (Dipl.-Phys.) in 2000. Since then he has worked in the Computer Graphics Group at University of Bonn, Germany. His work has focused on the measurement of surface reflection properties for computer graphics.

Reinhard Klein studied Mathematics and Physics at the University of Tübingen, Germany, from where he received his MS in Mathematics (Dipl.-Math.) in 1989 and his PhD in computer science in 1995. In 1999 he received an appointment as lecturer ("Habilitation") in computer science also from the University of Tübingen, with a thesis in computer graphics. Since October 2000 he is professor at the University of Bonn and director of the Institute of Computer Science II.



Figure 9. Comparison of a photograph (left) with a spectral rendering (center) and a RGB rendering (right) showing the accuracy of our data as well as the impacts of metamerism.

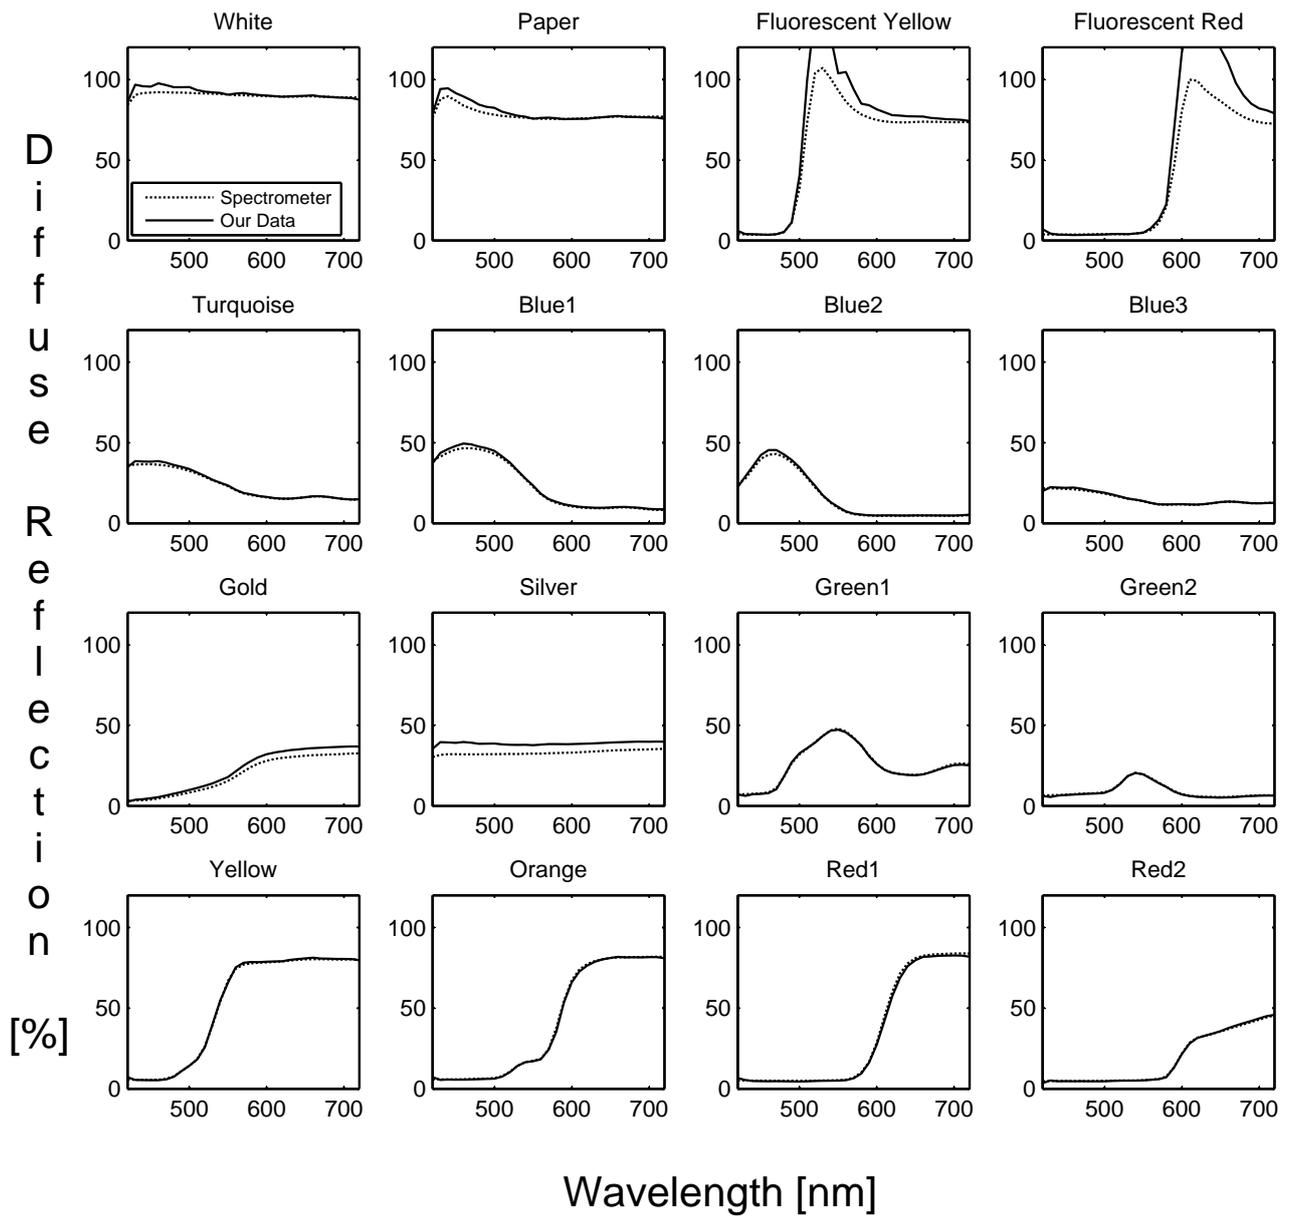


Figure 10. Comparison of our data integrated to mimic the spectrometer behaviour with measurements taken with the spectrometer. All fields but the specular and the fluorescent ones are reproduced very good.

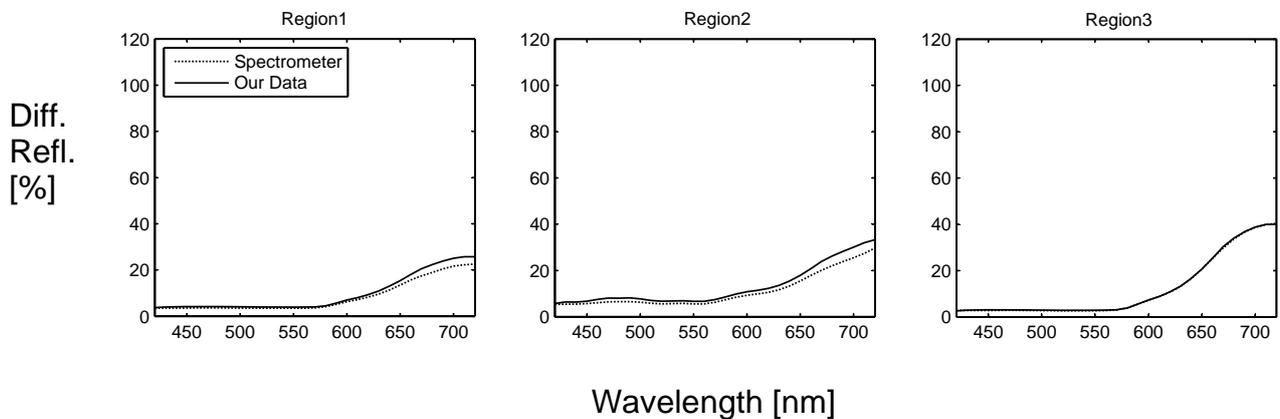


Figure 11. Comparison of our data integrated to mimic the spectrometer behaviour with measurements taken with the spectrometer. We tried to match the spots where we measured with the I1 as good as possible.