Supplemental Material for PaperID 0159 "A Reconfigurable Camera Add-On for High Dynamic Range, Multispectral, Polarization, and Light-Field Imaging"

This document provides an additional discussion of design choices and information regarding the processing steps involved in the applications.

1 Details on the Choice of the Diffuser

The diffuser structure is of major practical concern for image quality. Since we are using high resolution sensors, and essentially perform a 1:N minification, the diffuser structure should ideally be below $N \times \mu$, where μ is the physical pixel size, to avoid seeing the diffuser structure in the final image. A large diffusion angle, however, requires sufficient structures for scattering. This can either be achieved by larger surface structures (typically $50 - 100\,\mu m$) in relatively thin materials, or by finer scattering structures in thicker slabs. The former option leads to graininess in the image whereas the latter induces glare by multiple scattering. We experimented with both types of diffuser, one thin material that has been extracted from an LCD screen (thickness $\approx 0.1\,mm$), and one diffuser type that is designed for polarization-based 3D rear projection screens (thickness $\approx 1\,mm$ available from ScreenTech GmbH (material type "ST-Professional-DCF"). The latter diffuser is designed to be polarization-preserving. The two diffusers are at opposite ends of the spectrum, where the LCD diffuser shows considerable graininess, whereas the diffuser intended for 3D rear projection screens shows considerable glare. Fig. 1 shows the angular scattering profile of both diffusers and measurements of a target that is designed for the measuring of image resolution and contrast. The LCD diffuser shows graininess but much better contrast than the 3D rear-projection diffuser which suffers from glare. Unfortunately, the hotspot nature of the LCD diffuser prevents its use in our filter-based design since it passes light-field components into the different copies. This means that parallax effects are observable in the differently-filtered views – clearly an undesirable attribute. We therefore chose to work with the 3D rear-projection diffuser.

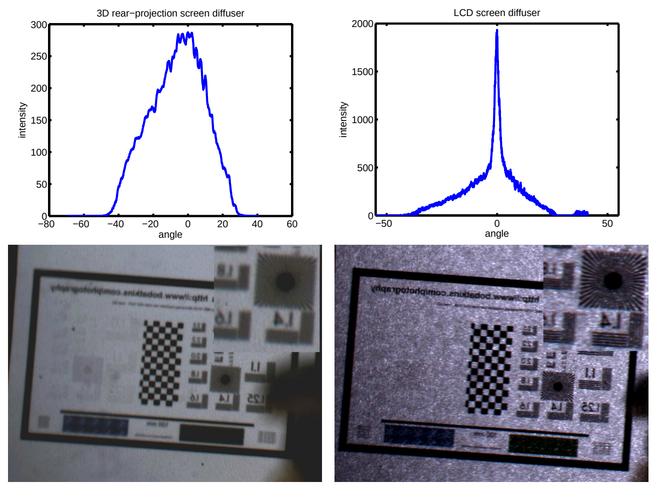

Figure 1: Top Row: *Angular response of two different diffusers.* Left: *A diffuser intended for 3D rear-projection. The scattering profile is broad and does not feature spiky peaks - it is said to be free of hot spots. This type of diffuser is usable in our system.* Right: *A diffuser extracted from an LCD screen. The diffuser profile shows a strong peak in the center - this peak is called a hot spot in diffuser optical language. It passes light-field components into our imaging system and is unsuitable for filter-based imaging.* Bottom Row: *Comparison of image quality.* Left: *Diffuser for 3D rear projection. The contrast is reduced due to glare within the thick material of the diffuser. In this example, the resolution is about 1.2 line pairs per mm.* Right: *LCD diffuser, the material shows surface structure which has a highly non-uniform angular response. The resolution and contrast, however, are much better than in the 3D rear-projection diffuser, 1.8 line pairs per mm can be resolved. Further, the LCD diffuser results in observable parallax between sub-images.*

2 Details on Multispectral Calibration and Estimation

2.1 Spectral Calibration

We took images of a MacBeth ColorChecker classic through all 116 broadband filters provided in the Roscolux swatchbook as produced by Rosco Laboratories. The transmission spectra of these filters where measured by a spectrometer (CCS 200, Thorlabs Inc.; $200\,nm - 1000\,nm$). As in the paper we denote them by $f_i, \ , i = 1 \dots 116$. In addition, we measured the spectral output of a light source that has a higher efficiency in the blue than an incandescent light bulb, a high pressure mercury vapor lamp s_{mv}. The 116 input images are radiometrically compensated. We then approximate the Bayer filter spectral transmission curves by a linear combination of basis functions.

$$f^{r|g|b}(\lambda) = \sum_{j=0}^{N-1} a_j^{r|g|b} \phi_j(\lambda). \tag{1}$$

Here, N is the number of basis functions ϕ_j used for the approximation and $a_j^{r|g|b}$ are the coefficients for the red, green, and blue response functions. Inserting into Eq. (3) from the paper, we obtain

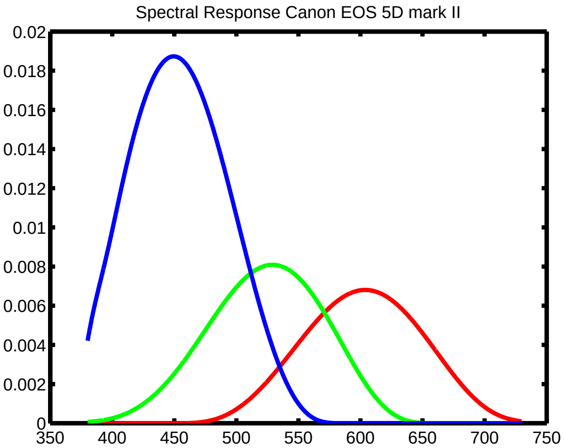

Figure 2: *Spectral response of the Canon EOS 5D mark II as estimated from* 116 *broadband filtered images.*

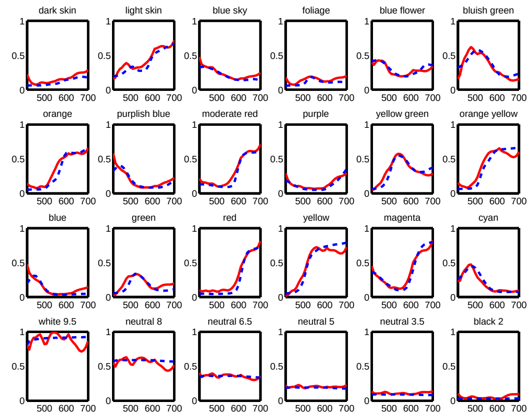

Figure 3: *Spectral reflectance of a MacBeth ColorChecker classic as imaged and reconstructed by our system (red) and average ground truth measurements obtained from 30 ColorCheckers (blue), source:* `http://www.babelcolor.com/main_level/ColorChecker.htm`.

a linear system of equations:

$$I_i^{r|g|b}(x,y) = \sum_{j=0}^{N-1} a_j^{r|g|b} \int_\lambda l_\lambda(x,y,\lambda), \phi_j(\lambda) f_i(\lambda) s_{mv}(\lambda) d\lambda, \tag{2}$$

with the integral now being completely determined. In theory, we could solve the resulting linear systems

$$\mathbf{I}^{r|g|b} = \mathbf{\Phi} \mathbf{a}^{r|g|b} \tag{3}$$

in a least-squares sense by inverting $\mathbf{a}^{r|g|b} = (\mathbf{\Phi}^T \mathbf{\Phi}) \mathbf{I}^{r|g|b}$ where the vectors collect the individual coefficients and image values and the matrix $\mathbf{\Phi}_{i,j}$ contains the integral values for the i'th image and the j'th basis function. In practice, we need to add a standard Laplacian regularization term (with Neuman boundary conditions) and enforce positiveness of the solution. We solve

$$\min_{\mathbf{a}^{r|g|b}} (\mathbf{\Phi} \mathbf{a}^{r|g|b} - \mathbf{I}^{r|g|b})^2 + \alpha (\mathbf{L} \mathbf{a}^{r|g|b})^2 \text{ subject to } \mathbf{a}^{r|g|b} \geq 0 \tag{4}$$

using the MATLAB `quadprog` sub-routine. In practice, we use 50 Gaussian basis functions to cover the range between $400\,nm$ and $700\,nm$, which is the region of best sensitivity of our camera. The regularization parameter α was set to 50. The resulting curves are shown in Fig. 2

2.2 Spectral Image Estimation

Our imaging device creates 9 copies that are filtered with spectral broadband filters selected from the Rosco Labs Roscolux swatchbook. The filters were selected manually, based on spectral coverage coniderations. The filters that were used for all experiments are { Cyan #4360, Yellow #4590, Red #26, Orange #23, Green #89, Blue-Green #93, Lavender #4960, Blue #80, Magenta #4760 }. The resulting images are RGB images, i.e. each of the nine filters is modulated by an additional Bayer filter, resulting in 27 measurement channels. We use the same algorithm as for the calibration of the spectral sensitivities, Sec. 2.1. There is one minor modification; following [Toyooka and Hayasaka 1997; Park et al. 2007], we use a PCA basis for the spectral dimension. It is well known, that natural spectral reflectances can be modelled by a low-dimensional linear model [Jaaskelainen et al. 1990]. This modification necessitates an adaptation of the Laplacian prior since regularization is to be performed on the reconstructed function, not on the coefficients of its linear expansion. For this purpose, an additional matrix $\mathbf{\Psi}$ containing the basis vectors in its columns is introduced into the regularization term.

$$\min_{\mathbf{a}^{r|g|b}} (\mathbf{\Phi} \mathbf{a}^{r|g|b} - \mathbf{I}^{r|g|b})^2 + \alpha (\mathbf{L} \mathbf{\Psi} \mathbf{a}^{r|g|b})^2. \tag{5}$$

Moreover, since a large number of pixels have to be reconstructed, we cannot afford to solve a quadratic program in this case. We therefore resort to solving the standard per-pixel least-squares problem

$$\mathbf{a}^{r|g|b} = (\mathbf{\Phi}^T \mathbf{\Phi} + \alpha (\mathbf{L}\mathbf{\Psi})^T (\mathbf{L}\mathbf{\Psi}))^{-1} \mathbf{I}^{r|g|b}. \tag{6}$$

We construct the PCA basis by analyzing a database of 1269 spectra measured from the Munsell book of colors (source: `http://www.uef.fi/spectral/munsell-colors-matt-spectrofotometer-measured`) in agreement with [Jaaskelainen et al. 1990]. In practice, we use the mean of the data set and the first 14 PCA vectors as basis functions. The regularization parameter α has been set to a low value of 0.01 since the basis itself has strong regularizing properties.

Evaluation To evaluate our reconstruction scheme, we recorded a MacBeth Color checker classic under the illumination of a high-pressure mercury vapor lamp. The results are compared to average measurements obtained from spectrometers that were collected on `http://www.babelcolor.com/main_level/ColorChecker.htm`, see Fig. 3. The results are in good agreement with the collected data.

2.3 Relighting and Simulation of Color-Deficiency

Our applications are based on a modulation of the acquired color spectra. The processing chain involves the computation of neutral spectral images with 29 wavebands. Since our examples have been recorded in a lab setting, the spectral composition of the illuminating source was known (measured by spectrometer Thorlabs CCS 200). We can then modulate the spectral image stack with illuminating spectra from different sources:

- a simulation of different black body spectra using Planck's law,

- an assortment of daylight spectra obtained from `http://www.uef.fi/spectral/daylight-spectra`,
- an assortment of artificial light spectra obtained from `http://www.uef.fi/spectral/artificial-lights`, and
- a simulated sweeping band pass with Gaussian profile of $20\,nm$ width.

The resulting spectral images were converted to XYZ using the CIE64 standard observer. Finally, a conversion to sRGB (gamma 2.2) with a simulated D65 daylight source was performed.

Simulation of Color Deficiency For simulating the vision of color deficient subjects, we obtained the color-matching functions for protanomalous (response of the red cones shifts towards green by about $10\,nm$) and deutanomalous subjects (response of the green cones shifts towards red by about $10\,nm$) [DeMarco et al. 1992]. These two are the most common color deficiencies encountered in the population. In addition, there are three types (protanopia, deutanopia, tritanopia) that have zero response for the L, M, and S cones, respectively. The simulations were performed by substituting the CIE64 color matching functions with their adapted versions. Since these are not designed to convert to XYZ but to an LMS cone response, we first mapped the output to XYZ, which was then converted to sRGB for display.

3 Detailed Derivation of the Horn-Schunck-based Depth Estimation

To derive a depth map, we do not rely on stereo matching, but instead employ a modified optical flow algorithm. We adapted the Horn and Schunck functional by introducing a scalar depth function $d(x, y)$ for the center view. This depth value results in a parallax displacement in the neighboring images via $d \cdot [u_i, v_i]^T$, where the vector $[u_i, v_i]^T$ denotes the direction of the epipolar lines in a neighboring view I_i. Our goal is to estimate a depth that best explains all views in a least-squares sense. The resulting error functional adapted from Horn & Schunck flow is given by

$$\min_d E = \int \sum_i (d(\triangledown I_i \cdot [u_i, v_i]^T) + I_i^t)^2 + \alpha^2 ||\triangledown d||^2 dxdy, \quad (7)$$

where i is an index denoting the neighboring views, excluding the center. $[u_i, v_i]$ are constant vectors that indicate the direction of the epipolar lines in view i, $\triangledown I_i = [I_i^x, I_i^y]$ is the spatial gradient of view i, and I_i^t denotes the "temporal" derivative between view i and the center view. The term $\alpha^2||\triangledown d||^2$ is a standard Laplacian smoothness term on the depth map with regularizing parameter α^2. The scalar function $d(x, y)$ is the quantity being optimized for. The corresponding Euler-Lagrange equation is

$$d \cdot \sum_i (I_i^x u_i + I_i^y v_i)^2 + \sum_i I_i^t (I_i^x u_i + I_i^y v_i) - \alpha^2 \triangle d = 0. \quad (8)$$

We solve it by standard techniques, discretizing the spatial derivatives by central differences with Neumann boundary conditions and the "temporal" derivative by Euler forward differencing. We use the Horn & Schunck approximation to the Laplacian $\triangle d = (d - \bar{d})$, where

$$\begin{aligned} \bar{d} = &\ 1/12 \quad (d_{x-1,y-1} + d_{x+1,y-1} + d_{x-1,y+1} + d_{x+1,y+1}) \\ + &\ 1/6 \quad (d_{x,y-1} + d_{x,y+1} + d_{x-1,y} + d_{x+1,y}) \,. \end{aligned}$$

The equation

$$d^{k+1} = \frac{\alpha^2 \bar{d}^k - \sum_i I_i^t (I_i^x u_i + I_i^y v_i)}{(I_i^x u_i + I_i^y v_i)^2 + \alpha^2}$$

then defines an update rule to solve for d by Jacobi iterations k. The scheme is implemented in a standard scale-space fashion [Meinhardt-Llopis and Prez 2012] to allow for large displacements.

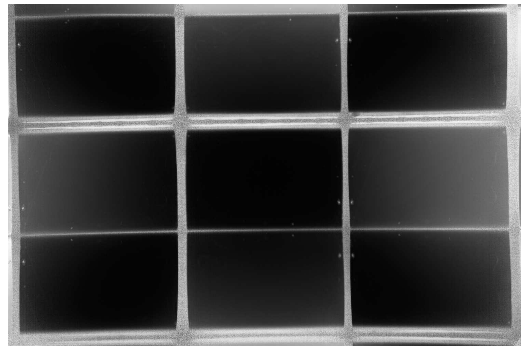

Figure 4: *Demonstration of how the optical system changes the polarization state of the incoming light. The image shows the distribution of degree of polarization in the camera's red color channel. The coding uses black to indicate totally polarized (100%) and white to show unpolarized (0%) states.*

4 Details on the Requirements and Calibration for Polarization Imaging

To allow for polarization imaging with our optical system, it is essential that we employ a polarization-preserving diffuser instead of a regular diffuser, since a regular diffuser essentially acts as a depolarizer. Light traversing through a depolarizer becomes unpolarized regardless of its initial polarization state. In other terms, the Mueller-matrix of a depolarizer is not a full row rank matrix and, therefore, retrieving the original Stokes vector of the light becomes impossible.

The other prerequisite is that we need to determine an effective Mueller-matrix $\mathbf{M_{sys}}(x, y)$ of the system for each pixel (x, y). It characterizes the optical system in the sense of how it changes the polarization state of the incoming light. In Fig. 4, we demonstrate this effect in an example image, which shows the measured linear degree of polarization p_{deg} across the nine sub-images in an optical system without any polarization filters. By definition, we can derive the degree of polarization $p_{deg} = \sqrt{({s^{(1)}}^2 + {s^{(2)}}^2)}/s^{(0)}$ from the Stokes vector $\mathbf{s} = [\, s^{(0)}\ s^{(1)}\ s^{(2)}\,]$. In this particular case, we let a diffused and completely polarized light beam through the setup. While at the center, the system retained the 100% degree of polarization, towards the edges the light became less polarized due to the reflections on the mirrors. We need to compensate for this change by carrying out a calibration procedure with the complete system using differently oriented polarizers in the filter array.

Basically, the calibration is a Mueller matrix polarimetry whose principles and possible realization scenarios are thoroughly described in [Goldstein 2003]. It needs a system with a complete polarization state generator (PSG) and a complete polarization state analyzer (PSA) component. For PSG, we have used an incandescent light source, a paper diffuser, and a polarizer that were placed at the entrance of the optical system. With these simple tools both states (totally linearly polarized and unpolarized) can be generated. We place these elements in the following order; for totally polarized light *illumination* $\rightarrow$ *paper diffuser* $\rightarrow$ *polarizer* $\rightarrow$ *optical system* and for unpolarized light *illumination* $\rightarrow$ *polarizer* $\rightarrow$ *optical system*. Here, we exploit the feature of the paper diffuser, which acts as an almost ideal depolarizer. Further, by still retaining the polarizer in the optical path, we can ensure that the light beam entering the optical system is going to have the same spectrum regardless of the generated polarization state. With various orientations of the polarizer, five different totally polarized states and one unpolarized state were reproducibly created. These states were measured by a ground-truth imaging polarimeter consisting of a camera, lens (the same as in our optical system) and a manually rotated polarizer (similar to [Neumann et al. 2008]). Ground-truth Stokes vectors $\mathbf{s}_{gt}^{(i)}$, $i = 1...6$, were determined by averaging the obtained Stokes-vectors within a window of 1500×1500 pixels in the center of the images, which were completely filled by the unfocused image of the diffused light beam. Then the same measurements of these states were carried out through the optical system, employing the same camera, lens, and the additional manually-rotated polarizer, yielding a Stokes vector $\mathbf{s}_{sys}^{(i)}(x, y)$ for each pixel (x, y). These Stokes-vectors are then matched against to the ground truth ones using the linear relation $\mathbf{s}_{sys}^{(i)}(x, y) = \mathbf{M}_{sys}(x, y)\mathbf{s}_{gt}^{(i)}$, $i = 1...6$. $\mathbf{M}_{sys}(x, y)$ is solved via a least-squares regression. Finally, the obtained system Mueller matrices $\mathbf{M}_{sys}(x, y)$ have to be registered with the same transformation that is applied to the sub-images and will yield five Mueller-matrices $\mathbf{M}_i$, $i = 1...5$ in each sub-image pixel. These matrices are used for retrieving pixel-by-pixel Stokes-vectors of the optical system itself and we can follow the description given in the main paper (see Eq.(4)).

References

DEMARCO, P., POKORNY, J., AND SMITH, V. C. 1992. Full-Spectrum Cone Sensitivity Functions for X-chromosome-linked Anomalous Trichromats. *JOSA 9*, 9, 1465–1476.

GOLDSTEIN, D. H. 2003. *Polarized Light, 2nd edition*. CRC Press, New York, US.

JAASKELAINEN, T., PARKKINEN, J., AND TOYOOKA, S. 1990. Vector-Subspace Model for Color Representation. *JOSA 7*, 4, 725–730.

MEINHARDT-LLOPIS, E., AND PREZ, J., 2012. Horn-Schunck optical flow with a multi-scale strategy. Preprint, Image Processing On Line. http://www.ipol.im/pub/pre/20/.

NEUMANN, L., HEGEDUS, R., HORVATH, G., AND GARCIA, R. 2008. Applications of High Precision Imaging Polarimetry. In *Proc. Computational Aesthetics in Graphics, Visualization and Imaging*, 89–97.

PARK, J.-I., LEE, M.-H., GROSSBERG, M. D., AND NAYAR, S. K. 2007. Multispectral Imaging Using Multiplexed Illumination. In *Proc. ICCV*, 1–8.

TOYOOKA, S., AND HAYASAKA, N. 1997. Two-Dimensional Spectral Analysis using Broad-Band Filters. *Optical Communications 137* (Apr), 22–26.