Practical Low Cost Recovery of Spectral Power Distributions

Sara Alvarez¹

¹Universidad de Zaragoza

Timo Kunkel² Belen Masia^{1,3}

²Dolby Laboratories

ries ³MPI Informatik

Abstract

Measuring the spectral power distribution of a light source, that is, the emission as a function of wavelength, typically requires the use of spectrophotometers or multispectral cameras. Here, we propose a low-cost system that enables the recovery of the visible light spectral signature of different types of light sources without requiring highly complex or specialized equipment and using just off-the-shelf, widely available components. To do this, a standard DSLR camera and a diffraction filter are used, sacrificing the spatial dimension for spectral resolution. We present here the image formation model and the calibration process necessary to recover the spectrum, including spectral calibration and amplitude recovery. We also assess the robustness of our method and perform a detailed analysis exploring the parameters influencing its accuracy. Further, we show applications of the system in image processing and rendering.

Categories and Subject Descriptors (according to ACM CCS): I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Radiometry

1. Introduction

Obtaining a multispectral image of a scene allows us to extract information of it that our eyes fail to capture. This has wide applicability in a variety of disciplines, including computer vision, medical imaging or computer graphics. Multispectral imaging enables characterization of light sources, material recognition and inspection, industrial control, spectral rendering, or tissue identification, among others. Besides, with the proliferation of cameras-both DSLRs and in mobile devices-, photographic hardware and software filters and simple image processing algorithms are widely used and available in a number of software programs and apps. These apps can benefit from knowledge of the spectral signature of the light source(s) illuminating a scene for different purposes: to yield better renditions of the images, for white balancing, for the application of certain artistic filters or for insertion of synthetic objects into the image. This spectral signature, also called spectral power distribution (SPD), is the power emitted by a light source as a function of wavelength, and can be measured with spectrophotometers. These devices, however, are specialized equipment of relatively high complexity and cost, not available to the wide public.

This gives rise to this paper, where we propose a system to estimate the SPD of different light sources at a low

submitted to COMPUTER GRAPHICS Forum (6/2015).



Figure 1: Left: Elements of our capture system, which employs a DSLR camera and a diffraction filter like the ones shown. Right: Two images of light sources captured with our camera using a 500 lines/mm diffraction filter, showing the diffraction pattern (spectral signature) of the source.

cost and requiring only off-the-shelf equipment. This equipment comprises a camera and a diffraction filter, such as those shown in Figure 1, *left*, and described in Section 3. Our main focus is simplicity of use, availability, and low cost. Other approaches in the literature have used prisms for a similar purpose [DTCL09, CDT*11], but require a larger setup, while we just require a filter placed in front of the lens. Further, we do not aim at obtaining an emission spec-

trum as accurate as that yielded by spectrophotometers, but to obtain a spectrum that can be used for a number of commonplace everyday applications, such as white balancing or compositing of synthetic objects into an image (Section 7). This paper is an extension of the method previously published in [APKM14]. We extend that work by performing a detailed analysis in order to test the robustness of our approach. In this analysis, we explore how different variations in the scenario or the light source affect the quality of the recovered emission spectrum. We also compare ourselves to the work of Cao et al. [CDT*11]: While we obtain slightly less precise results, our setup is less bulky and complex and requires less equipment, and still enables useful applications.

2. Related Work

In 1814 Joseph von Fraunhofer became known as the discoverer of the missing dark absorption lines in the solar spectrum [KB32], an indicative of the relationship between the emission spectrum and the chemical composition of the molecules of the emitter. Years later, in 1898, J. J. Thomson would develop the first mass spectrometry method [Tho21], which was later further developed and improved by F. W. Aston [Ast19] and A. J. Dempster [Dem18]. One of the main goals of both mass spectrometers and spectrophotometers is to measure the spectral power distribution of a signal, typically by splitting the ray into its wavelengths. To achieve this, dispersive or diffractive elements must be placed in the light path. These elements are usually gratings, where the splitted rays are spatially shifted according to the grating equation [Hec87].

This paper focuses on spectral measurement techniques based on a spatial multiplexing of the spectrum. The reader may refer to compilations such as [IWH10] for a thorough description of spectral acquisition methods; we summarize here some of the main approaches. These techniques usually use dispersive or diffractive elements together with optical elements to change the rays direction in a scene that will impinge the camera sensor creating multiple, spectrally sampled images which will be reconstructed in post-processing. An example are methods based on computed tomographic imaging spectrometry (CTIS) [DD95, VMHD07]; they were developed by Okamoto and Yamaguchi [OY91] using diffraction gratings, a field stop and a lens to estimate a 3D distribution from a set of 2D images.

To capture specific spectral components of a scene, several methods based on placing band-pass filters in front of the camera's sensor have been designed. Examples include tunable filters (ELTs) [Gat00], or a wheel of filters [WH04]. Typically, these methods employ narrow-band filters which give rise to a low light throughput; to avoid this, Toyooka and Hayasaka [TH97] present a system based on broadband filters together with computational inversion and without any scanning mechanism. A high number of approaches to tunable filter systems have also been made: Liquid Crys-

tal Tunable Filters (LCTFs) [Cha10], based on a cascade connection of Lyot filter stages; variable optoacoustic tunable filters, where a crystal was acoustically excited working as a diffraction grating; or an association of redirecting mirrors [GKT09, GFHH10, HFHG*05] which achieves images of up to 25 spectral bands in real-time. Bodkin et al. [BSN*09] and Du et al. [DTCL09, CDT*11] exploit a similar concept by using a set of pinholes to restrict the rays captured by the sensor. Behind these pinholes a prism is located for spreading the spectrum, which then impacts in the sensor, allowing for the acquisition of multispectral videos. Other approaches to multispectral imaging include the use of coded illumination [PLGN07], or the use of one color channel to increase the spectral resolution [KN07], trading off temporal for spatial resolution. In comparison with the above, the system we present in this paper minimizes the complexity of the hardware setup and the calibration process. Note that, while these works are the closest to ours, in this paper we do not aim at recovering a multispectral image or video of a scene (i.e., spectrally sampled values per pixel); instead, our goal is to obtain the SPD of a light source in a simple, practical manner using a low cost, commonly available setup.

3. Overview of the System

We adopt a computational photography approach, in which we will trade off the spatial dimension for the spectral one. Our system requires only a diffraction filter mounted on a camera; this can be done simply using a filter mount, which is screwed in at the end of the lens. This filter will separate light emitted by the source into its spectral components, producing a diffraction pattern (spectral signature) that we will capture and process to obtain the spectral power distribution of the source. Figure 1, right, shows two examples of diffraction patterns captured with our Canon EOS 500D and a linear diffraction filter. Diffraction filters such as the ones used in this work (see Figure 1, left) are easily found and common among photographers and amateur astronomers, and they can be bought at a cost lower than five US dollars. The image formation model of the system is described in Section 4.

Obtaining the spectral power distribution of the source from the diffraction pattern implies inverting a number of processes that the signal undergoes in the camera. We refer to this as the calibration of the system. Its first step is the socalled *spectral calibration* (Section 5.1), where the mapping between sensor pixels of the spectral signature and their corresponding wavelengths is found. The second step deals with *amplitude recovery* (Section 5.2), since the camera has varying sensitivities for each wavelength and color channel that need to be inverted to recover the original signal. While we do provide a mathematical model for these different aspects of the system, there are a number of issues which are difficult to take into account theoretically. These issues include technical specifications of cameras (e.g., spectral sensitivity curves) not being disclosed by manufacturers, or influencing phenomena that need very complex modeling (e.g., glare or sensor bleeding). We thus revert to a data-driven approach in some cases to circumvent this.

4. Image Formation Model

Our capture system, described in Section 3, allows us to capture an image of the light source with an associated diffraction pattern (or spectral signature) such as the ones shown in Figure 1, *right*. Next, a 1D intensity profile is obtained from a scanline of this captured image; this scanline is taken from the central line of the spectral signature, and although it is currently done manually by simply selecting a line, automatization is straightforward. This section describes the image formation model of the system, while the next delves on how to transform the aforementioned 1D intensity profile in an actual spectral power distribution. We describe here the



Figure 2: Ray diagram of the system. Light rays are diffracted by the filter, stopped down by the aperture, and impinge on the sensor. For simplicity, a pinhole aperture is represented instead of the lens system. Note that here two diffraction orders (m = -1 and m = 1) are shown; in practice, the position of the light source is offset in the sensor, so that only one diffraction order (m = 1) is visible.

model for a 1D linear diffraction filter, but other filters could be used; a ray diagram of the setup can be seen in Figure 2. For simplicity, in the following derivations we will assume a thin lens model. As we will see later in Section 5.1 and Figure 5, *right*, this model proves to be enough for our purposes, since we use this derivation mainly to support the validity of our data-driven approach.

A light source imaged through the linear filter gives rise to two symmetrical diffraction patterns at a distance from the central maximum. Diffraction in the filter is governed by the diffraction grating equation, such that the diffracted angle β_j at which the peak occurs for a certain wavelength λ_j is given by the following equation:

$$\beta_j = \arcsin(\frac{m\lambda_j}{d} - \sin\theta) \qquad \qquad j = \{1, 2, \dots, N\}, \quad (1)$$

where m is the diffraction order, d the separation between grooves of the grating, θ the angle of incidence to the grating, and N is the number of samples considered along the spectrum. The light source is assumed to be far away from the filter, such that $\theta \approx 0$. Details on the derivation of this equation can be found, e.g., in Hecht's book [Hec87]. We work in the Fraunhofer regime, as given by the Fresnel number $F = a^2/(L\lambda)$, a being the width of the slit or groove (50% of the separation between slits), and L the distance between the grating and the diffraction plane (the sensor). The separation between lines or grooves (d) can be obtained from the specifications of the grating (in our case, 500 lines/mm). We work with the first order of diffraction, i. e., m = 1, since the second and higher orders are much lower in magnitude and fall outside the sensor area for typical diffraction gratings. Note that this means that the field of view when taking the image needs to be such that both the light source and the first-order diffraction pattern need to be present in the captured image. Diffraction patterns for a point and a finite fluorescence light source captured with a linear diffraction filter can be seen in Figure 1, right.

We can compute the displacement in the sensor from the central maximum (corresponding to m = 0) for each wavelength λ_j as:

$$x_j = s' tan(\beta_j)$$
 $j = \{1, 2, ..., N\},$ (2)

where the diffraction angle β_j for each λ_j is given by Equation 1. Equations 1 and 2, together with the well-known thin lens equation 1/f = 1/s + 1/s' (Gaussian Lens formula [Hec87], where *f* is the lens' focal length, *s'* the virtual distance of the image and *s* the distance between the aperture plane and the light source), and knowing the effective pixel size *t*, allows us to compute the relationship between distance in sensor pixels measured from the central maximum $(m = 0) p_j$ and wavelength λ_j (Figure 2):

$$p_j = \frac{sf}{t(s-f)} tan(arcsin(\frac{m\lambda_j}{d} - sin\theta)) \quad j = \{1, 2, ..., N\}$$
(3)

This relationship will be used for the spectral calibration described in Section 5.1. Besides, we can compute the spectral resolution (*res*, nm/pixel) of our system by obtaining the position in the sensor for the extremes of the visible spectrum:

$$res = \frac{\lambda_M - \lambda_m}{p_M - p_m},\tag{4}$$

where subindices M and m refer to the maximum and minimum wavelengths in the visible spectrum, and p_M and p_m

submitted to COMPUTER GRAPHICS Forum (6/2015).



Figure 3: Diffraction pattern from a white low-energy consumption bulb after deconvolution using the Richardson-Lucy algorithm.

can be obtained from Equation 3. In our system, this yields a resolution of 0.333 nm/pixel.

4.1. Finite light sources

With the exception of ideal point light sources, the shape of the light bulb modulates the emission spectrum, effectively mixing different wavelengths of the spectral signature. This modulation can be modeled as a 2D convolution between the shape of the bulb and the image of the spectrum, since it acts as part of the point spread function of the system [Goo05]. To cast aside the influence of the light bulb shape we apply a Richardson-Lucy deconvolution [Ric72, Luc74]. The kernel for this deconvolution is obtained by cropping and thresholding the image. This operation can be done automatically by detecting the maximum value of the image (which corresponds to the light source) and setting a relative threshold around it. Figure 3 shows the image of a source and its spectrum before and after deconvolution. We chose Lucy-Richardson for deconvolution because it is ubiquitous, with plenty of implementations available, and yields considerably less noisy results in our scenario than other standard deconvolution methods. This is shown in Figure 4, which depicts the spectrally calibrated distributions obtained with different common deconvolution methods for a white fluorescent light source. Richardson-Lucy deconvolution assumes a Poisson distribution of the noise of the sought image; however, in our setup, our distribution would be better approximated by a combination of Gaussian and Poisson distributions, for which other more sophisticated methods could be better suited [GRPMSM11].

4.2. Background subtraction

There might be cases in which our light source of interest is not the only source of illumination present in the scene. In those cases, our system requires two photographs, one with the diffraction filter and one without the filter, as input. Light



Figure 4: Spectrally calibrated intensity distributions from a white low-energy consumption light source obtained using different deconvolution methods. In reading order: Wiener deconvolution (with a NSR of 10); Wiener deconvolution with a regularizer consisting of a Laplacian operator; deconvolution relying on a prior enforcing sparsity of gradients (designed for natural images) [LFDF07]; and Richardson-Lucy deconvolution, which provides a smoother, less noisy and overall better solution, thus being the selected method.

is additive and follows the principle of superposition, meaning we can simply subtract both images to obtain a dark image with the spectral signature of the source of interest. The diffraction undergone in the filter, however, is a non-linear process, and therefore this subtraction will not only leave behind the diffraction pattern of the light source that we want, but also residual light due to diffraction of light coming from parts of the scene other than the light source. This stray light, however, is orders of magnitude lower than the diffraction pattern of interest (recall that we are using HDR images, allowing for more dynamic range), and thus has no significant effect on its recovery. We further discuss this issue and show the validity of this approach in Section 8. Following the subtraction, we perform a deconvolution as described in Section 4.1 (this step is not necessary if we are dealing with a point light source) and proceed with the calibration as described in Section 5.

5. Calibration of the System

This section covers the process to transform the 1D profile extracted from the HDR image of the light source captured with the diffraction filter into a spectral power distribution (SPD) which can be used for different applications, such as the ones shown in Section 7. This is what we refer to as calibration of the system. Although not strictly necessary, we work with HDR images; this offers us a higher dynamic range and avoids the need to linearize the intensities in the image. Further, HDR capture is commonplace and comes as a preset mode in a large number of cameras today.

Several factors come into play when attempting to retrieve the emission spectrum from an image. The diffraction undergone in the filter (recall that diffraction is a non-linear process), the presence of noise, dispersion and glare as light propagates within the camera, human factors present during capture, or lens aberrations are all responsible of modifications performed to the original emission spectrum. The lack of detailed specifications of the camera's hardware and software due to manufacturers' reticence to disclosure further hinders the generation of a comprehensive forward model and thus of a reliable inverse model. Here, we focus only on the effects that the camera's hardware plus the filter cause on the captured signal. These effects must be reversed to obtain the original emitted signal as accurately as possible.

The calibration consists of two parts: first, a *spectral calibration*, in which the objective is to find the mapping function between pixel values and light wavelength values; second, *amplitude recovery* in which we model variations in the amplitude of the SPDs due to the camera's varying sensitivity to different wavelengths.

5.1. Spectral Calibration

The mapping between the displacement in pixels within the sensor (from the central maximum) and the wavelength is given, theoretically, by Equation 3. We are, however, making a series of assumptions, such as the use of a thin lens model, when deriving that equation. Therefore, we additionally take a practical approach by imaging a light source with known SPD, and obtaining the displacement (in image pixels) of a series of characteristic peaks of the source (Figure 5, left) for which the corresponding wavelength is known. These corresponding points (shown as black circumferences in Figure 5, right) allow us to fit an experimental mapping using least squares. We consequently obtain the linear relationship shown in Figure 5, right. Note that the theoretical curve (Equation 3) already hinted this linear relationship (for incidence angle $\theta \approx 0$). We use R^2 as a measure of the goodness of fit of this mapping, which yields $R^2 = 0.99$. We also plot in Figure 5, right, the theoretical mapping for the parameters of our system (f = 50mm, s = 1.97m, m = 1, $\theta \approx 0$, $d = 2.004 \ \mu m, t = 9.38 \ \mu m/pixel$). There is a slight difference between both mappings, probably due to the simplified modeling of the lens system we perform. The fact that the relationship very well approximates a linear function means that two recognizable peaks of any light source will suffice to calibrate such a system. This is very practical, since common fluorescent light sources have well-known peaks (546.5 and 611.6 nm for white fluorescent light sources [DTCL09]) which can be used for spectral calibration without the need of knowing the whole SPD of the source used for calibration or the specific parameters of the system implementation. Thus, we will use the experimental calibration.



Figure 5: Left: Emission spectrum of a white fluorescent light source exhibiting several characteristic peaks at known wavelengths, and used for the experimental calibration. Right: Mappings pixel-wavelength obtained via the theoretical (red) (Equation 3) and the experimental (blue) spectral calibration procedures. The points used for the experimental mapping are shown as black circumferences.



Figure 6: Red curves correspond to the 1D profiles of the diffracted intensities emitted by a white fluorescent point light source (left) and a white LED point light source (right) after spectral calibration (see text for details). The corresponding ground truth SPDs are shown in black for comparison purposes.

This spectral calibration is applied to 1D intensity profiles obtained from the captured images. The spectral calibration allows us to map the pixel positions of the profile into wavelengths, obtaining the profiles shown in Figure 6 for two different light sources. Once the system is spectrally calibrated, the next step is the recovery of the correct amplitude values of the SPD, described next.

5.2. Amplitude Recovery

A digital camera's sensor is overlaid with a color filter array that causes the impinging light to be weighted by different wavelength-dependent response curves corresponding to the spectral sensitivities of the three color channels (see Figure 7, left).

In an ideal case, Equation 5 would be fulfilled for each color channel $i = \{R, G, B\}$ for each pixel in the 1D profile. We will denote each pixel *p* by its corresponding wavelength λ obtained by means of the spectral calibration.

$$r(i,\lambda) = c(i,\lambda)s(\lambda), \tag{5}$$

where $r(i, \lambda)$ is the relative intensity extracted from the im-

submitted to COMPUTER GRAPHICS Forum (6/2015).



Figure 7: Left: Canon EOS 500D normalized spectral sensitivities $c(i,\lambda)$ for $i \in \{R,G,B\}$ [BCC*09]. Right: Weighting values ω as a function of $v(i,\lambda)$, the normalized distance to the median of the three color channels for quotient $r_{i,\lambda}/c_{i,\lambda}$ for a certain λ (see text and Equation 8 for details).

age for that wavelength and color channel, $s(\lambda)$ the emission of the light source for a certain wavelength, and $c(i,\lambda)$ the spectral sensitivity of the camera for each color channel *i* and wavelength λ .

Following Equation 5, recovering the light source's emitted spectrum would imply a simple division for each pixel and color channel, the three color channels yielding the same $s(\lambda)$ value for a certain pixel. This ideal situation is however not the case in our real scenario, where bleeding among pixels and other factors introduce inaccuracies. We therefore solve the overdetermined problem that Equation 5 constitutes via a least squares optimization. Besides, the values of $c(i, \lambda)$ that fall below a certain threshold ($c(i, \lambda) < 0.02$) have been ignored in the optimization because of their very low signal-to-noise ratio. We look for the estimation of the emitted spectrum, $\hat{s}(\lambda)$, that minimizes the following energy term E:

$$E = E_{data} + w_s E_{smooth} \tag{6}$$

The first term, E_{data} , corresponds to the error with respect to the observed data, while the second term, E_{smooth} , is a regularization term that strives for smoothness and continuity. Both are detailed below. Finally, w_s controls the relative weight of each term; $w_s = 50$ works well across sources and is used for all the results shown in the paper.

The term E_{data} is constructed from Equation 5 as follows (note that we use $c_{i,\lambda}$ for $c(i,\lambda)$, and equivalently for *r* and \hat{s} , for clarity of the notation):

$$E_{data} = \sum_{i=1}^{3} \sum_{\lambda} \left(\omega_{i,\lambda} \left(\hat{s}_{\lambda} - \frac{r_{i,\lambda}}{c_{i,\lambda}} \right) \right)^2, \tag{7}$$

where, again, *i* corresponds to the color channel. The data term thus tries to minimize the error with respect to the observed data, looking for the \hat{s}_{λ} that best approximate the $r_{i,\lambda}/c_{i,\lambda}$, but this error is weighted by values $\omega_{i,\lambda}$. This weighting function ω penalizes the errors of data points that yield values of the quotient $r_{i,\lambda}/c_{i,\lambda}$ very far away from its median for the three color channels for that λ , assuming they

are more unreliable. The function, shown in Figure 7, *right*, is given by the following expression:

$$\omega_{i,\lambda} = 1 - exp\left(\frac{-1}{\nu(i,\lambda)^2 + \alpha}\right) \cdot exp\left(\frac{1}{1 + \alpha}\right), \qquad (8)$$

where values $v(i, \lambda)$ are the normalized distances of the data point to the median of the corresponding three color channels, given by:

$$v(i,\lambda) = \frac{1}{P_i} \frac{r_{i,\lambda}}{c_{i,\lambda}} - median\left(\left\{\frac{1}{P_k} \frac{r_{k,\lambda}}{c_{k,\lambda}}\right\}_{k=1}^3\right)$$
(9)

$$P_i = \max_{400 \le \lambda \le 700} \left\{ \frac{r_{i,\lambda}}{c_{i,\lambda}} \right\} \qquad \forall i = \{R, G, B\}$$
(10)

Parameter α controls the slope of the function ω , and it takes a value of $\alpha = 0.09$ for fluorescent light sources and of $\alpha = 10$ for incandescent light sources, whose spectral power distributions typically are more broadband and vary more smoothly than those of fluorescent sources (see also Section 6.1 for further discussion on this).

The smoothness term for the obtention of \hat{s} is given by:

$$E_{smooth} = \sum_{\lambda} (\Delta \hat{s}_{\lambda})^2, \qquad (11)$$

where the discrete Laplacian is implemented as a circular convolution with the 1D kernel $\begin{bmatrix} 1 & -2 & 1 \end{bmatrix}$.

The minimization of *E* as given by Equations 6 to 11 allows us to recover the original amplitude of the spectrally calibrated signal, as shown in Figure 8, *right*. Note that λ in Equations 7 to 11 is in practice a discrete signal λ_j with $j = \{1, 2, ..., N\}$, as in Section 4; the subindex *j* has been dropped here for clarity. More results of recovery are shown in Section 6.



Figure 8: The red curve corresponds to the 1D profile of the diffracted intensity emitted by a white incandescent point light source after spectral calibration (left). The blue curve shows the estimated SPD of this light source after recovering the amplitude through the algorithm we have exposed (right). The corresponding ground truth SPD is shown in black for comparison purposes.

6. Results and Analysis

The technique presented allows us to recover the spectral signature of both point and finite light sources from high dynamic range images captured with the camera and diffraction filter setup. To quantify the accuracy of our estimations, we captured ground truth SPDs of the recovered light sources with an Ocean Optics USB4000-UV-VIS photospectrometer, which has a spectral resolution of 1.5 nm in the range 200-850 nm. The sensitivity of the Canon 500D is limited to the range between 400 and ca. 680 nm (see spectral sensitivity curves in Figure 7, *left*), and thus so is our recovered spectrum. This range, however, covers the majority of the visible spectrum and suffices for a large number of applications (human spectral sensitivity as given by the CIE 1924 photopic luminous efficiency function is below 1.7% of its maximum value for $\lambda > 680$ nm [RKAJ08]). In this section we provide additional results, offer an extensive analysis of the performance of our system under different conditions, and compare our technique with the state of the art.

6.1. Precision analysis

One of the main advantages of our method lies in recovering a sufficiently accurate spectrum even in largely underconstrained scenarios. To prove this robustness, we have performed a thorough analysis testing the accuracy of the recovered spectra under different conditions. This analysis also allows us to validate the steps in our pipeline, including the deconvolution step (Section 4.1) and the background subtraction (Section 4.2). In particular, we test the influence of three factors, shown in Figure 9. The first factor is the shape of the source, that is, whether the light source is a point or a finite one. This tests whether the deconvolution required for finite light sources has a negative influence on the final result. Second, we investigate if the background has an effect on the recovered spectra. This factor has three levels: (i) dark background, in which the background holds no information (it is black), and the recovered source is the only source of light; (ii) non-dark background, in which the background has some kind of color information, which gets mixed with the spectral signature (and will be removed by background subtraction); and (iii) illuminated background, in which there are additional sources of illumination in the scene which may interfere in the measurement. Testing this factor is also a means of validating the background subtraction step of our pipeline. Last, we investigate whether the type of light source has an effect on the quality of the recovered spectrum. To do this, we classify our sources into two large types: incandescent and fluorescent, which are the two main types of artificial sources to be found in the real world.

Our dataset consists of six different light sources (white, green and red low energy consumption fluorescent sources, and white, green and red incandescent sources), each captured with all the different backgrounds and both as a finite and as a point light source (the latter we obtain by placing a black screen with a hole in front of the light source). This yields a total of 36 images (3 backgrounds \cdot 2 shapes \cdot 6 light sources = 36) for which we recover the SPD. We measure the accuracy

Factor	F-statistic	p-value
shape	3.17	0.088
background	1.65	0.213
type	26.33	0.000

Table 1: ANOVA results when testing the influence of the three factors (shape, background and type) in the error of the recovered SPD. A p-value below 0.05 for a factor indicates that the null hypothesis (the factor does not influence the error) can be rejected for the corresponding factor.

of our recovered spectra by computing the mean squared error (MSE) with respect to the measured ground truth. Some of these light sources, together with their recovered spectra and associated ground truth can be seen in Figure 11, where we plot the average of the six measurements (3 *backgrounds* \cdot 2 *shapes*) taken per light source. The main features of the light sources are captured, even if there are inaccuracies in the recovered spectrum. These recovered curves are enough to allow for light source recognition, which can be used to enable a number of applications, explained in Section 7.

We analyze the errors in the recovered spectra using an ANOVA with three factors (background, shape, type). The goal is to know which of the factors have a significant effect on the error of the recovered spectra. The results of the ANOVA are shown in Table 1. Neither the background nor the shape of the light source have a significant effect on the error (p - value > 0.05 for both factors), showing that the background subtraction and the deconvolution steps work well in practice. The type of light source does have a significant effect on the recovered spectra (p - value = 0.000). These results are further shown in Figure 10, in which we plot the average error for each level of each factor. As expected, the error in dark backgrounds is slightly lower than in the other cases, but the difference is not significant. In the case of shape, point light sources have slightly larger error (although not significant) than finite light sources. This is probably because only finite light sources are deconvolved, and point light sources are not purely a point (an impulse) but actually have a finite diameter. Further, the point light sources have a lower signal-to-noise ratio (the signal intensity is lower), which may increase the error in the recovery. Finally, incandescent light sources have significantly larger error than fluorescent ones. This is due to the characteristics of their spectra, which are typically broadband, and have a large amount of energy in areas in which the spectral sensitivity of the camera is relatively low. This is especially true in the higher wavelengths ($\lambda \ge 670nm$), were incandescent sources emit a large part of their energy, while the sensitivity of the camera is very poor. Meanwhile, in fluorescent light sources, due to their narrowband nature around certain wavelengths, there tends to be a color channel with high sensitivity around those wavelengths which is able to provide a

submitted to COMPUTER GRAPHICS Forum (6/2015).



Figure 11: Average spectral power distributions recovered with our method (blue curves) and corresponding ground truths (black curves) for four different light sources used in our experiments. From left to right: white, red and green low-energy consumption (fluorescent) and red incandescent sources. Insets show the spectral signature of the sources.

good SNR for recovery. It is due to these characteristics of the spectra of incandescent sources that in our amplitude recovery step we use an α value (see Equation 8) of 10, higher than that of fluorescent sources: In the latter, due to the aforementioned narrowband nature, there tends to be a color channel that is significantly more reliable than the other two, and a low value of α (large slope in the ω weights function given by Equation 8) favors this selective reliability of color channels. Meanwhile, in incandescent sources, the reliability of all color channels is more similar, which is taken into consideration by means of a higher α . Additionally, incandescent light sources have a smoother spectral power distribution, and larger α values increase the weight of the smoothness term with respect to the data term. To further illustrate the performance of our system when recovering a variety of incandescent light sources, we show in Figure 12 the recovered spectra for a green and a yellow incandescent light source together with their ground truth; the graphs show how the main features of the light sources are recovered.

6.2. Comparison to other methods

One of the main advantages of our method in contrast with existing methods is that it does not require any specialized equipment, or any complex or bulky setup. At the same time, note that the goal of most previous works [CDT*11] is multispectral imaging (i.e., obtaining an image of a scene separated in spectral components), while our goal is to recover the spectral emission of a light source of the scene. In this section, we aim at showing that, for this purpose, our method can obtain comparable results to those of more complex techniques with simpler equipment. In particular, we compare ourselves to the results obtained by the work of Cao et al. [CDT*11]. In their paper, the authors present a system for capturing multispectral videos. The system consists of a lightproof box containing a prism, a monochrome camera, an occlusion mask and an optional optical filter. They also provide a result of light spectrum recovery for a fluorescent light source (Figure 12(b) in their paper), to which we compare.

Figure 13 shows the spectral emission of a fluorescent



Figure 9: Images illustrating the factors (and their corresponding levels) taken into account in our analysis of performance. From top to bottom: shape of the light source, type of light source, and background. We perform an ANOVA in order to assess which of these factors, if any, has an influence on the recovered SPD.



Figure 10: Average MSE (mean squared error) for each level of each of the three factors (shape, background and type). Bars show the standard error of the mean.

submitted to COMPUTER GRAPHICS Forum (6/2015).



Figure 12: Recovered (blue curves) and ground truth (black curves) SPDs for a green incandescent finite light source (left), and a yellow incandescent finite light source (right). In both cases, the main characteristics of the light sources i.e., the central peak and subsequent decay in the green incandescent source, and the lack of emission until 500 nm followed by a monotonic increase in the yellow source—are captured by the system.

light source recovered with Cao et al.'s method and with our method (blue curves). In both cases the corresponding ground truth is shown (black curves). The figure also shows the mean squared error (MSE) in both cases. To compute it for the case of Cao et al., we have sampled the curves every 5 nm. In our case, we show our most difficult case: a finite light source captured with a non-dark, illuminated background (i.e., there was an additional, incandescent light source in the scene). While our system is slightly less precise ($MSE_{Cao} = 0.006$; $MSE_{ours} = 0.015$), even in the worst case scenario it offers enough accuracy to enable many applications, and is simpler. Note that their system requires, in addition to a prism, a mask and a more bulky setup. Other multispectral imaging methods [GKT09, GFHH10, HFHG*05, PLGN07] are not comparable to our system because they require a much more complex setup (e.g., cascaded birefringent interferometers in [GFHH10] or coded illumination in [PLGN07]), and their aim is different: They recover spectral values per pixel of a scene, while we recover the SPD of a light source.

6.3. Recovery in the wild

The analysis described in Section 6.1 already shows that our system is robust to different backgrounds and shapes of light sources. To further demonstrate this capability, we include here an example of a highly challenging scenario, showing the results obtained in this case and their accuracy. Figure 14 shows, for two different light sources, a photo of the scenario (top row), a photo of the scenario captured with the diffraction filter (middle row), and the recovered spectrum (bottom row). We can see how, for both sources, our system is able to recover the main features of the spectral power distributions of the light sources, despite the complex shape of the bulb (and thus of the photographed spectral signature) and the background information. Further, we provide white bal-



Figure 13: Left: SPD obtained with the method by Cao et al. [CDT* 11] for a fluorescent light source (blue curve) with corresponding ground truth (black curve). They show the recovered SPD for different widths of the slits of their mask, but the 0.2 mm width (blue curve) is the best case (figure adapted from [CDT* 11]). Right: SPD obtained with our method for a white fluorescent light source (blue curve) with corresponding ground truth (black curve). We slightly sacrifice accuracy for a much simpler, easily available setup, and can still recover the light source's main features.

ancing results for this scene captured with an incandescent source in Figure 16.

7. Applications

7.1. Light Source Recognition

Spectral emission of commercial light sources are available in a number of databases [KGH*14, Sed09]. In our case, we have built up our own database by capturing ten common light sources, including e.g., LEDs, a laser, low-energy consumption bulbs, or incandescent sources.

Given a recovered SPD, which can be seen as a vector \hat{s}_{λ} , it is compared to the distributions in the database by simply computing the Euclidean distance (L2-norm) between both vectors, once they are appropriately re-sampled for equal dimensionality. Other distances such as the Fréchet distance [CCdVE*10] were tested but yielded worse results. This recognition step has been used in the two following applications shown in this section.

7.2. White Balancing

There is a large number of white balancing algorithms [Lam05, HCWW06, BGS07]. In most cases, however, they rely on obtaining information of the illuminant either from the image once captured [HMP*08, CFB02], or from very simple calculations done by the camera at the time of capture [Shi06, II02]. We show here how to perform white balancing via recovery of the SPD of the illuminant, yielding appealing results in scenarios where other existing methods would fail.

We first compute the CIE 1931 XYZ color space values [RKAJ08] from the SPD of the light source under consideration. This XYZ values are computed for the spectral

submitted to COMPUTER GRAPHICS Forum (6/2015).

S. Alvarez, T. Kunkel, and B. Masia / Practical Recovery of Spectral Power Distributions



Figure 15: Three scenarios processed with four white balancing algorithms. From left to right: original images, white balanced images using the light source's correlated color temperature algorithm (our method), the Gray World algorithm, the Max-RGB algorithm, and the Shades-of-Gray algorithm [Cohl1]. These scenarios are illuminated by different kinds of sources; from top to bottom: filament incandescent, opaline and filament incandescent together, and a yellow incandescent light source. Note that we provide results at least on par with the state of the art and consistent across a variety of illuminations.



Figure 14: Top row: photographs of the scenario. Middle row: photographs of the scenario with the diffraction filter added to the lens. Bottom row: SPD of the light source recovered with our method (blue curve) and corresponding ground truth (black curve). The left column corresponds to a white fluorescent light source and the right column to a red incandescent source. The original images are high dynamic range and have been tone-mapped for display.

emission stored in the database for the light source identified via the recognition system (Section 7.1) getting as input the spectrum recovered as explained in Section 5. This is done to increase the accuracy of the XYZ values computation, since the recovered spectrum accumulates error when integrating over the spectrum to obtain the XYZ values.

Once the XYZ values are calculated from the SPD, a white balancing algorithm is used, based on finding out the correlated color temperature *CCT* of the light source which illuminates the scene [McC92, HALR99]. To compute the *CCT*, first the *x* and *y* chromaticity coordinates are calculated from the XYZ values [RKAJ08], and then McCamy's polynomial equation for *CCT* is used [McC92], as follows:

$$n = \frac{x - 0,3320}{y - 0,1858} \tag{12}$$

$$CCT = -449n^3 + 3525n^2 - 6823, 3n + 5520, 33$$
(13)

The *CCT* is measured in Kelvin and gives good results within the range (2000 - 12500K). This value shall be converted to sRGB values, for which the D_{65} illuminant's white point \mathbf{e}_w is used. To this end, we use a conversion table [Cha01] which gives us the sRGB values e_i (i = {R,G,B}) for the *CCT* of the light source. These values e_i are used to compute a weighting value g_i per color channel *i* of the image we want to white balance (*I*), such that the white balanced image I^{WB} is given by $I_i^{WB} = g_i I_i$. The weighting values g_i are given by:

$$g_i = \frac{e_{w_i}}{e_i} \qquad \forall i = \{R, G, B\}, \tag{14}$$

submitted to COMPUTER GRAPHICS Forum (6/2015).

where \mathbf{e} is the vector defined by the three sRGB values of the light source under consideration.

Some results of our white balancing algorithm in three different scenarios can be seen in Figure 15. We also show a comparison against three widely used white balancing algorithms, described in [Coh11]. These three algorithms are based on image post-processing, estimating the illuminant from image pixel values and modifying the pixel values of each RGB channel accordingly. Some current digital cameras perform white balancing via some calculations at the time of capture [Shi06, II02]; we do not compare against them, however, because the algorithms are typically proprietary information, not disclosed by manufacturers, and they operate at intermediate stages performed in the camera firmware which we do not have access to.

The fact that our white balancing algorithm does not depend on the spatial distribution of the pixels in the scene yields an overall better performance than the others, which may fail in specific cases, as shown in Figure 15. The image may not satisfy the Gray World Assumption (which states that given an image with enough variety of RGB colors, the average values of these three color channels should average out to a common gray value), as is the case in Figure 15, first row; or be illuminated with a skewed SPD that would disrupt the behaviour of the Shades-of-Gray algorithm (Figure 15, third row). A common example of images in which image-based white balancing algorithms such as the above will fail is that of images with a clearly dominant color or range of colors. Finally, Figure 15, second row, also shows that our algorithm can be applied when there is more than one light source illuminating a scene, yielding compelling results. Compared to existing approaches, we provide results at least on par with them and consistent over a range of illuminations. We have further applied this white balancing to the challenging case presented in Section 6.3 and show in Figure 16 both the scene illuminated by an incandescent source (left) and the white balancing result (right) when employing the emission spectra of the identified incandescent source. Even though the scene is very challenging (see examples of spectral signatures in that scene in Figure 14), our method is able to obtain convincing white balancing results.



Figure 16: Original photograph (left) and white-balanced result (right) of a scene illuminated with an incandescent light source. Insets show close-ups of the scene where the white balancing is better appreciated (see white text).

submitted to COMPUTER GRAPHICS Forum (6/2015).

7.3. Compositing of Synthetic Objects

Integrating synthetic objects in a convincing way within a captured image of a real scene is always a challenging task, requiring manual work to painstakingly adjust the lighting. If, however, the spectral emission of the light source illuminating the scene is available, which can easily be done with our system, compositing becomes easier and offers more compelling results in less time.

The SPD obtained with our system is once more used for identification of a light source, whose spectrum is input to a spectral rendering engine to render an image of the object which we want to integrate in an image (in this case, the Stanford bunny). During rendering, rays will be characterized by a certain frequency (wavelength); in our case, the implementation utilizes 20 nm steps for the sampling in wavelength. Once the object is rendered, it is introduced in the captured image by using as a mask an alpha channel created when rendering the 3D model. Figure 17 shows some results for the same scene illuminated under two dramatically different light sources; note that the bunny's reflectance is white and purely diffuse.



Figure 17: Spectrally-rendered Stanford bunny inserted in a scene illuminated with a green low-energy consumption light (top row) and a yellow incandescent (bottom row).

8. Discussion and Limitations

An error in the estimation of the SPD will always exist, since part of the hardware of the camera is not known. It is not our goal here to surpass the accuracy of a spectrophotometer, but to offer a simple approach to obtain an estimation of the SPD of a light source that is good enough for applications in imaging that can benefit from having this spectral information. In addition to unknown hardware specifications (including the camera's spectral sensitivities, which are not provided by the manufacturer but obtained from other sources instead), there are a number of more complex phenomena, such as bleeding of pixels into neighboring pixels in the sensor, glare, or noise, which are very hard to model theoretically; this is one more reason for our experimental approach to calibration. We use for recovery an HDR image of the light source, since this allows us to avoid possible saturation in the sensor, and yields a linear relationship between impinging radiance and intensities recorded in the image. This, together with the fact that HDR creation software and capture capabilities are becoming commonplace, makes them a better choice than LDR images. Currently, the 1D profile is extracted from this HDR image manually by clicking on the start and end points of the spectrum in the image. This is a very simple requirement from the user; however, if desired, the process could be automatized by, e.g., looking for the row with maximum integrated intensity.

In the presence of multiple light sources or a non-dark background, background subtraction allows us to extract a spectral signature for the recovery, as explained in Section 4.2. Having an HDR image as input allows the possible stray light arising from diffraction in other sources of light (or diffuse surfaces) other than our source of interest to be negligible. This can be seen in Figure 10, in which the graph shows the error in the recovery of light sources captured with different backgrounds: The error, although slightly smaller, is not significantly different in scenes with a dark background (no stray light), and scenes with a non-dark background or with an illuminated background (requiring subtraction), such as that shown in Figure 9 (bottom right). A further example is provided by the scenario in Figures 14 and 16, which has two light sources (the one shown plus another white incandescent source not directly visible to the camera), and a white diffuse background that could contribute stray light to the diffraction pattern. For six different light sources (white, red and green incandescent and fluorescent sources), we computed the contrast ratio between each sample (pixel) in the diffraction pattern and its surrounding background, in the image resulting after background subtraction and deconvolution. This ratio was on average well over two orders of magnitude (grand average across light sources yielded a ratio of 585). Thus, background subtraction provides a reasonable approximation even in the presence of other sources of light in the scene that may give rise to diffracted stray light. An exception to this would be a scene with abundance of highly reflective surfaces, but that is typically not the case.

Additionally, if there are multiple light sources in the scene, they can be individually recovered using our procedure. There would be, however, a limitation to recovery, if the spectral patterns are not isolated and those of two or more light sources interfere. While modeling and inverting this interference could be possible, it could also result in an ill-posed problem, and the accuracy could be compromised.

Other diffraction filters, such as cross filters, could also be used. The use of these filters introduces redundancy in the data, since a number of diffraction patterns appear in the image, and this can make the estimation more robust. The downside, however, is the presence of interference which needs to be computed and inverted, and as before, could result in an ill-posed problem. Further, the spectral calibration curve (Figure 5, *right*) will no longer yield a linear relationship, which significantly eases calibration allowing to perform it with any common white fluorescent source.

As mentioned before, in terms of wavelength range we are limited to the part of the spectrum that our sensor is sensitive to, that is, from 400 to ca. 700 nm. Additionally, the different sensitivities to different wavelengths (Figure 7, *left*) cause the signal-to-noise ratio to be different at different wavelengths, and thus not all wavelengths are recovered with the same accuracy. An option would be to introduce this factor into the optimization described in Section 5.2. The effect of these possible inaccuracies is variable with the light source under consideration (e.g., underestimating the main peak of a fluorescent has a larger effect that underestimating that same wavelength in an incandescent source); this can have an important effect in applications such as white balancing, and is the reason why we now include an identification and ground truth retrieval step prior to the actual white balancing. It should also be noted that proper amplitude recovery depends on an accurate spectral calibration, since otherwise the spectral sensitivity curves are not adequately matched to their corresponding wavelengths.

Finally, light sources often have an anisotropic emission, that is, their emitted spectrum is different for different angles, since the emitted intensity varies. While we do not measure angular variation here, nothing prevents the usage of the system taking images from different directions to characterize angular variation. In any case, the SPD would typically vary very slightly with the angle, if at all.

9. Conclusions and Future Work

We have proposed a low-cost system for acquisition of spectral power distributions of light sources using off-theshelf equipment consisting of a photographic camera and a diffraction filter. We present a calibration method so that a 1D intensity profile extracted from an HDR image of the light source is transformed into the SPD emitted by the light source. This calibration implies obtaining the mapping between pixels and wavelengths, which as we show can be done by imaging a common white fluorescent, whose two main peaks are well-known. The second step recovers the original amplitude of the emitted spectrum, and requires knowing the spectral sensitivity curves of the camera. We have performed a series of tests to assess the influence of a number of factors in the recovered spectra, and determine that the recovery is robust to different scene backgrounds and light source shapes. We have further shown a number of applications in image processing which can benefit of knowing the SPDs of light sources illuminating a scene. Even in highly challenging scenarios, our recovery is enough to enable these applications, such as white balancing. Note that the goal of our method is to provide SPDs that are accurate *enough* for these applications, while requiring easily available, low cost hardware.

As future work, and as hinted in Section 8, information from the camera's spectral sensitivities could be employed when recovering the amplitude of the emitted spectrum as an indication of the reliability of the estimation obtained from the data for a particular color channel and wavelength. Further, in the deconvolution step, other more sophisticated algorithms that take into account both Poisson and Gaussian noise, such as that of Gil-Rodrigo et al. [GRPMSM11], could be used. Another interesting option would be to combine our approach with perceptual aspects taking into account the characteristics of the human visual system (e.g., [MMG11, MPCG12]).

Potentially, obtaining the spectral signature can be done not from direct imaging of light sources, but of other objects. This, which could be challenging for a diffuse object, may be worth looking into for specular or glossy objects. Further, this approach could be extended to obtaining spectral information from materials, aimed at, for instance, material recognition [DTCL09].

Acknowledgements

We thank Lara Presa and Diego Gutierrez for their valuable insights and input, and Adrian Jarabo and Adolfo Muñoz, for fruitful discussions and help in the rendering part. We would also like to thank the reviewers for their insightful comments. This research has been funded by the Spanish Ministry of Economy and Competitivity under project LIGHT-SLICE (TIN2013-41857-P). The work has also been partially supported by the Max Planck Center for Visual Computing and Communication.

References

- [APKM14] ALVAREZ S., PRESA L., KUNKEL T., MASIA B.: Low cost recovery of spectral power distributions. In *Proceedings of CEIG* (2014). 2
- [Ast19] ASTON F. W.: LXXIV. A positive ray spectrograph. The London, Edinburgh, and Dublin Phil. Magazine and J. of Science 38, 228 (1919), 707–714. 2
- [BCC*09] BOSWELL B., CARROLL P., CARTER K., CHEN J., NEWTON D., ET AL.: DxOMark Labs. Industry standard for camera and lens independent image quality measurements and ratings. http://dxomark.com/Reviews/Canon-500D-Tli-vs. -Nikon-D5000/Color-blindness-sensor-quality/, 2009. 6
- [BGS07] BIANCO S., GASPARINI F., SCHETTINI R.: Combining strategies for white balance. In *Electronic Imaging* (2007), pp. 65020D–65020D. 9
- [BSN*09] BODKIN A., SHEINIS A., NORTON A., DALY J., BEAVEN S., WEINHEIMER J.: Snapshot hyperspectral imaging: the hyperpixel array camera. In SPIE Defense, Security, and Sensing (2009), pp. 73340H–73340H. 2
- [CCdVE*10] CHAMBERS E. W., COLIN DE VERDIÈRE É., ER-ICKSON J., LAZARD S., LAZARUS F., THITE S.: Homotopic Fréchet distance between curves. *Computational Geometry* 43, 3 (2010), 295–311. 9

submitted to COMPUTER GRAPHICS Forum (6/2015).

- [CDT*11] CAO X., DU H., TONG X., DAI Q., LIN S.: A prismmask system for multispectral video acquisition. *Pattern Analy*sis and Machine Intelligence, IEEE Transactions on 33, 12 (Dec 2011), 2423–2435. 1, 2, 8, 9
- [CFB02] CARDEI V. C., FUNT B., BARNARD K.: Estimating the scene illumination chromaticity by using a neural network. JOSA A. 19, 12 (2002), 2374–2386. 9
- [Cha01] CHARITY M.: Software development. Blackbody color datafile. http://vendian.org/mncharity/dir3/blackbody/ UnstableURLs/bbr_color.html, June 2001. 10
- [Cha10] CHANNEL SYSTEMS INC.: Liquid crystal tunable filters. http://www.channelsystems.ca/ SpectralImaging-varispec-tunable.cfm, 2010. 2
- [Coh11] COHEN N.: A color balancing algorithm for cameras. EE368 Digital Image Processing (2011). 10, 11
- [DD95] DESCOUR M., DERENIAK E.: Computed-tomography imaging spectrometer: experimental calibration and reconstruction results. *Applied Optics* 34, 22 (1995), 4817–4826. 2
- [Dem18] DEMPSTER A.: A new method of positive ray analysis. *Physical Review Vol. 11*, N. 4 (1918), 316–325. 2
- [DTCL09] DU H., TONG X., CAO X., LIN S.: A prism-based system for multispectral video acquisition. In *Proc. of ICCV* (2009), pp. 175–182. 1, 2, 5, 13
- [Gat00] GAT N.: Imaging spectroscopy using tunable filters: a review. In *AeroSense* (2000), pp. 50–64. 2
- [GFHH10] GORMAN A., FLETCHER-HOLMES D., HARVEY A.: Generalization of the Lyot filter and its application to snapshot spectral imaging. *Opt. Express* 18, 6 (2010), 5602–08. 2, 9
- [GKT09] GAO L., KESTER R. T., TKACZYK T. S.: Compact image slicing spectrometer (ISS) for hyperspectral fluorescence microscopy. *Optics express Vol. 17*, N. 15 (2009). 2, 9
- [Goo05] GOODMAN J. W.: Introduction to Fourier Optics 3rd Ed. Roberts & Company, 2005. 4
- [GRPMSM11] GIL-RODRIGO E., PORTILLA J., MIRAUT D., SUÁREZ-MESA R.: Efficient joint poisson-gauss restoration using multi-frame L2-relaxed-L0 analysis-based sparsity. In Proc. of ICIP (2011), pp. 1385–1388. 4, 13
- [HALR99] HERNÁNDEZ-ANDRÉS J., LEE R. L., ROMERO J.: Calculating correlated color temperatures across the entire gamut of daylight and skylight chromaticities. *Applied Optics* 38, 27 (1999), 5703–5709. 10
- [HCWW06] HUO J.-Y., CHANG Y.-L., WANG J., WEI X.-X.: Robust automatic white balance algorithm using gray color points in images. *IEEE Trans. on Consumer Electronics* 52, 2 (2006), 541–546. 9
- [Hec87] HECHT E.: Optics. Addison-Wesley, 1987. 2, 3
- [HFHG*05] HARVEY A. R., FLETCHER-HOLMES D. W., GOR-MAN A., ALTENBACH K., ET AL.: Spectral imaging in a snapshot. In *Biomedical Optics* (2005), pp. 110–119. 2, 9
- [HMP*08] HSU E., MERTENS T., PARIS S., AVIDAN S., DU-RAND F.: Light mixture estimation for spatially varying white balance. ACM Trans. on Graph. 27, 3 (2008), 70. 9
- [II02] IDE M., ISHIMARU T.: Camera capable of white balance correction, 2002. US Patent 7,006,135 B2. 9, 11
- [IWH10] IHRKE I., WETZSTEIN G., HEIDRICH W.: A theory of plenoptic multiplexing. In CVPR (2010), pp. 483–490. 2
- [KB32] KORFF S., BREIT G.: Optical dispersion. Reviews of Modern Physics 4, 3 (1932), 471. 2

- [KGH*14] KOROW D., GOTTI M. B., HOWLEY J., HUG-GLER G., RAHM S., SIERLEJA R., SEDLAK S.: GE Lighting. Commercial products & solutions. Spectral power distribution curves. http://gelighting.com/na/business_lighting/ spectral_power_distribution_curves/, 2014. 9
- [KN07] KIDONO K., NINOMIYA Y.: Visibility estimation under night-time conditions using a multiband camera. In *Intelligent Vehicles Symp.* (2007), pp. 1013–1018. 2
- [Lam05] LAM E. Y.: Combining gray world and retinex theory for automatic white balance in digital photography. In *Proc. of ISCE* (2005), pp. 134–139. 9
- [LFDF07] LEVIN A., FERGUS R., DURAND F., FREEMAN B.: Image and depth from a conventional camera with a coded aperture. ACM Transactions on Graphics 26, 3 (2007), 70:1–70:9. 4
- [Luc74] LUCY L. B.: An iterative technique for the rectification of observed distributions. *The Astronomical J.* 79 (1974), 745. 4
- [McC92] MCCAMY C. S.: Correlated color temperature as an explicit function of chromaticity coordinates. *Color Research & Application 17*, 2 (1992), 142–144. 10
- [MMG11] MCNAMARA A., MANIA K., GUTIERREZ D.: Perception in graphics, visualization, virtual environments and animation. In *In SIGGRAPH Asia Courses* (2011). 13
- [MPCG12] MASIA B., PRESA L., CORRALES A., GUTIERREZ D.: Perceptually-Optimized Coded Apertures for Defocus Deblurring. *Computer Graphics Forum 31*, 6 (2012). 13
- [OY91] OKAMOTO T., YAMAGUCHI I.: Simultaneous acquisition of spectral image information. Optics letters 16 (1991), 1277–1279. 2
- [PLGN07] PARK J.-I., LEE M.-H., GROSSBERG M. D., NAYAR S. K.: Multispectral imaging using multiplexed illumination. In *Proc. of ICCV* (2007), pp. 1–8. 2, 9
- [Ric72] RICHARDSON W. H.: Bayesian-based iterative method of image restoration. J. Optical Society of America (JOSA) 62, 1 (Jan 1972), 55–59. 4
- [RKAJ08] REINHARD E., KHAN E. A., AKYZ A. O., JOHNSON G. M.: Color imaging: fundamentals and applications. AK Peters, Ltd., 2008. 7, 9, 10
- [Sed09] SEDRICK J.: Light database comparison. http:// zenware.com/light-database-comparison/, 2009. 9
- [Shi06] SHIMADA Y.: Electronic camera which detects flash influence on an image and controls white balance in accordance with the flash influence, 2006. US Patent 7,148,922. 9, 11
- [TH97] TOYOOKA S., HAYASAKA N.: Two-dimensional spectral analysis using broad-band filters. *Optics communications 137*, 1 (1997), 22–26. 2
- [Tho21] THOMSON J. J.: Rays of positive electricity. Cambridge Univ. Press, Cambridge (1921), 84. 2
- [VMHD07] VANDERVLUGT C., MASTERSON H., HAGEN N., DERENIAK E. L.: Reconfigurable liquid crystal dispersing element for a computed tomography imaging spectrometer. In *Defense and Security Symposium* (2007). 2
- [WH04] WANG S., HEIDRICH W.: The design of an inexpensive very high resolution scan camera system. In *Computer Graphics Forum* (2004), vol. 23, Wiley Online Library, pp. 441–450. 2

14