

DIY hyperspectral imaging via polarization-induced spectral filters

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

Abstract—We propose a novel design for a do-it-yourself hyperspectral imaging system which operates by taking multiple photographs through tunable, polarization-induced, spectral filters. Prior approaches in this do-it-yourself arena achieve hyperspectral imaging by selecting from a *discrete* set of spectra baked into existing products. In contrast, our approach is capable of generating a *continuous* family of broadband transmission spectra by simple rotations of stacked polarizers and waveplates. This greatly expands the potential range of representable spectra from a fixed-dimensional to an arbitrary-dimensional space. We analyze the theoretical spectral gamut of our approach and demonstrate its viability for spectral surface reflectance reconstruction both in simulation and with a low-cost physical prototype. Our prototype demonstrates that our approach can achieve comparable quality to prior work at reduced cost, while the new design space holds ample opportunity for increased quality and flexibility with professional manufacturing.

Index Terms—hyperspectral, polarized light, birefringence



1 INTRODUCTION

EVERY color in the vivid medley that surround us is composed of a spectrum of wavelengths. Digital cameras (and our eyes) project this continuous spectrum down to only three values, roughly corresponding to short (blue), medium (green), and long (red) wavelengths, throwing away most of the spectral information. While digital cameras are designed to replicate human perception faithfully, which make them well suited to capturing the aesthetic qualities of a scene, this information loss results in metamerism – infinitely many differing spectra that are nevertheless perceptually indistinguishable from one another.

Instruments capable of attaining high-fidelity spectral information are typically expensive, costing thousands to tens of thousands of dollars, and accessible only to research or industrial institutions. Note that we include in this category not only instruments billed as “hyperspectral cameras,” but also a range of other optical instruments such as fluorimeters and spectrophotometers. Such instruments are often used to study medical tissue samples and microorganisms or cells treated with fluorescent dyes, for example. Spectral information is useful, however, even in a variety of settings outside of well-funded research laboratories, such as tracking the aging of wines and beers, dating historic paintings, spotting harmful algal blooms in waterways, or determining the composition of stars or nebulae through telescopes. Collected spectral data could also be used as input to spectral renderers [1], [2] to faithfully simulate realistic images.

Many recent papers have taken up the gauntlet of do-it-yourself spectral reconstruction [3]–[10]. A common approach to this problem is to observe the spectral response of a subject under several *spectrally varying environments*, e.g. observing the subject under illuminants with different emission spectra or sensors with different spectral response curves. These methods have typically been limited to choos-

ing spectral variants from *discrete* sets of fixed spectra. In this work, we propose a novel design for do-it-yourself spectral reconstruction via spectrally tunable filters comprised of polarizers and waveplates that are capable of creating a *continuum* of transmission spectra, greatly expanding the space of spectrally varying environments that can be produced. While many have noticed the mesmerizing aesthetic effect these filters are capable of producing [11], this phenomenon has also been used in other scientific contexts, for example tracking the position and orientation of objects in an indoor scene [12]. While our theory could apply to polarizers and waveplates at any price range, we realize our design on a portable, low-cost prototype constructed from only an ordinary digital camera, lens filters obtainable at any casual photography store, and clear packing tape. We explore the capabilities of this prototype to acquire spectral surface material properties using spatial or temporal multiplexing of our filters (or both), and demonstrate that it can achieve comparable quality to prior work at reduced cost.

2 RELATED WORK

Commercial hyperspectral cameras (such as Olympus’s FV3000 Confocal Laser Scanning Microscope or Thorlabs’ Hyperspectral Imaging Cerna Rig) typically acquire a series of monochromatic images using a series of narrow bandpass filters (less than 10 nm) in order to capture every spectral band in the visible wavelength range individually. These setups typically cost thousands of dollars and have a microscope-like capture stage, restricting the subjects of interest to be small and flat. Alternate approaches to hyperspectral imaging typically leverage wavelength-dependent effects such as dispersion, diffraction, or birefringence in order to acquire spectral information.

Dispersion-based approaches

Coded aperture single-shot imaging (CASSI) techniques [13]–[16] place a coded binary mask in front of a dispersive

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

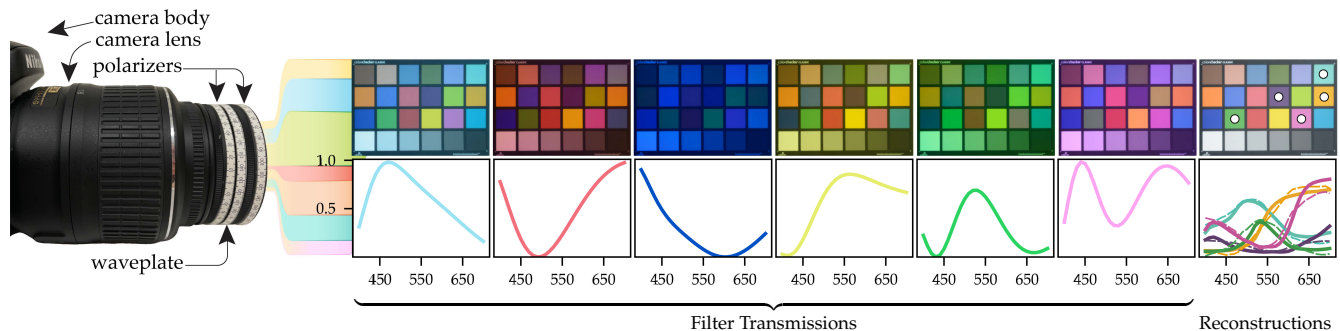


Fig. 1: Our filters, built from a stack of waveplates and polarizers, create a vivid array of transmission spectra that allow us to reconstruct the spectra of incoming light (*far right*: solid: ground truth, dashed: our reconstructions).

prism, spreading the spectral information over a small spatial neighborhood on the sensor. Du et al. [17] found that using even a simple checkerboard mask could recover spectral information. These setups have typically been confined to the laboratory due to the need for precise calibration and fabrication of the tiny, high-resolution aperture mask.

Notably, in the DIY space, Baek et al. [3] recover spectral information by taking a single photograph through a glass prism in front of an ordinary digital camera. Their approach harnesses cues from color fringing at the edges of objects in the scene in order to reconstruct the incoming spectra. While compact and inexpensive, their approach relies on having sufficient sharp edges in the scene and struggles with spiky spectra. Similarly, Hostettler et al. [18] leverage dispersion to project (not capture) color images by filtering white light through glass prisms and masks printed on transparencies.

Diffraction-based approaches

Computed tomography techniques use a diffraction grating to transform spectral into spatial information [19], [20]. In order to combat the spatial-spectral resolution tradeoff, Saragadam and Sankaranarayana [21] use an adaptive approach that alternates between acquiring spectral and spatial information using a diffraction grating and a coded aperture. Habel et al. [4] create a simplified diffraction-based hyperspectral camera by attaching a PVC pipe outfitted with several lenses and a diffraction grating to the front lens of an ordinary digital camera. Their approach, however, requires many scientific optical elements and a lengthy calibration process with extra equipment. Jeon et al. [22] simplify further by placing one custom diffractive element on top of a bare sensor, but fabricating the diffractive element requires photolithography and reactive-ion etching on a silica wafer. Monakhova et al. [23] replace the diffractive element with a tiny spectral filter array on the bare sensor, which also must be fabricated in a lab.

Birefringence-based approaches

Lyot filters use stacks of birefringent crystals in order to create narrow-bandpass filters, often used for tuning lasers. However, the shape of the transmission spectrum of these filters is fixed at manufacture time, although the transmission peak location may be shifted by rotation or voltage application. In contrast, our method creates transmission spectra that span the full visible spectrum but are tunable in shape.

In a similar vein to dispersion- or diffraction-based approaches, birefringence can also be used to split incoming light into diverging ordinary and extraordinary rays [24]–[27]. However, we harness the polarization-altering properties of the waveplates to create spectral filters, not to split light into diverging rays for spatial multiplexing.

In concurrent work, Sankaranarayanan et al. [28] harness the voltage-modulated birefringence of liquid crystals in order to create a programmable filter array that can be placed on top of a digital sensor. Their approach can be seen as a variation of our own that is able to achieve high accuracy via more expensive hardware that requires more engineering expertise to build and more extensive calibration.

Upsampling from RGB

Others have mapped RGB values to spectra directly by fitting box functions [29]–[31], basis functions derived from a large database [32], or smooth parametric models [33], [34]. Unlike our method, these approaches do not attempt to capture the spectral properties of a specific scene, rather they propose a general mapping from RGB to spectra.

Deep-learning-based approaches

Like many image processing tasks, deep learning has also shown promise for hyperspectral imaging, either by learning RGB upsampling from examples [35]–[39] or by fusing diffusion- or diffraction-based approaches with deep learning (see [40] for a recent survey). However, hyperspectral image datasets are still relatively small and sparse and may not apply to scientific settings. Notably, in the DIY space, Zhao et al. [5] train a neural network to reconstruct spectra from images taken through a mask of randomly distributed ink droplets printed by an ordinary consumer printer.

Linear-reconstruction-based approaches

Most related to our approach is the cornucopia of do-it-yourself, linear-reconstruction-based hyperspectral imaging systems that have been proposed in recent years. These typically take several photographic measurements of a scene while modulating the spectra of illuminants [6]–[8], filters over an illuminant [9], or sensor responses [10], [41]. The overarching theory behind these approaches is to observe the response of the spectra of interest under different *spectral environments*, and then piece together those observations into a linear system and solve for the unknown spectra.

While we adopt the same mathematical framework as these prior works, our approach differs in several key ways. The above methods are limited to choosing from a discrete set of possible spectral environments, whereas ours allows for continuous spectral variation through the design of our filters (Sec. 4). We also do not share some of the limitations specific to other design spaces, for example we are not limited to indoor spaces, unlike a system with controlled illuminants, nor do we need to compute homographies for image registration, as one would for a system with multiple sensors. With our spectrally varying filters built from polarizers and waveplates, taking a measurement is as easy as rotating the optical elements. Lastly, our prototype of the proposed method is especially affordable, using only components that can be acquired at any casual photography store, and does not require calibration with expensive laboratory equipment.

3 IMAGE FORMATION MODEL

(Throughout this work, we denote scalars and functions in lowercase (x), vectors in bold lowercase (\mathbf{x}), and matrices in bold uppercase (\mathbf{X})). A digital camera assigns an intensity to color channel k at pixel p by integrating the scene’s light spectrum with the channel’s spectral response curve $c_k(\lambda)$ over all visible wavelengths Λ :

$$p_k = \int_{\Lambda} c_k(\lambda) r(\lambda) e(\lambda) d\lambda. \quad (1)$$

Here, we have separated the scene into an illuminant spectrum $e(\lambda)$ and the surface reflectance spectrum $r(\lambda)$ of an object. If a filter with transmission spectrum $t(\lambda)$ is placed in front of the camera lens (or illuminant), the image formation model becomes:

$$p_k = \int_{\Lambda} c_k(\lambda) t(\lambda) r(\lambda) e(\lambda) d\lambda. \quad (2)$$

All linear-reconstruction-based approaches share this fundamental mathematical form, which we can abstract further by condensing the above quantities into an unknown spectrum $u(\lambda)$ and a (known) measurement spectrum $m(\lambda)$:

$$p_k = \int_{\Lambda} m(\lambda) u(\lambda) d\lambda = \langle m, u \rangle. \quad (3)$$

Typically $m(\lambda) = c_k(\lambda) t(\lambda) e(\lambda)$ and reflectance $r(\lambda)$ is the unknown $u(\lambda)$ we wish to solve for. However, $u(\lambda)$ could be any of the individual quantities above (or a combination). Likewise, any or all of the components $c_k(\lambda)$, $t(\lambda)$, or $e(\lambda)$ could be modulated in order to vary the spectral environment and provide further information toward recovering $u(\lambda)$ (see Table 1 for a breakdown of $m(\lambda)$ versus $u(\lambda)$ for prior work). On the right-hand side of Eq. (3), we replace the integral with inner product notation $\langle \cdot, \cdot \rangle$. This could be the inner product integral across wavelengths or a more general inner product that first represents the spectra in a lower-dimensional basis space. Given M distinct measurement spectra¹ $m_1(\lambda)$, $m_2(\lambda)$, \dots , $m_M(\lambda)$ and corresponding pixel measurements p_1 , p_2 , \dots , p_M , the

1. All approaches that use RGB cameras inherently vary $c_k(\lambda)$, producing three measurements with three separate color channels in one capture.

mentioned linear-reconstruction-based methods try to minimize the following loss:

$$\mathcal{L} = \arg \min_u \sum_{i=1}^M (p_i - \langle m_i, u \rangle)^2. \quad (4)$$

We propose a novel approach toward deducing the unknown spectrum whose key strength is the creation of a continuous family of measurement spectra $m(\lambda)$ of arbitrarily large dimensionality (Sec. 4). Given this continuous design space, we then address how we choose discrete sets of $m(\lambda)$ in practice (Sec. 5.1). Once a set of measurement spectra $m_1(\lambda)$, $m_2(\lambda)$, \dots , $m_M(\lambda)$ are chosen and associated pixel measurements p_1 , p_2 , \dots , p_M acquired, we solve Eq. (4) by assembling them into a linear system and finding a least-squares solution to the unknown spectra $u(\lambda)$ (Sec. 5.2). Lastly, we discuss applications of our method to surface reflectance (Sec. 6).

4 FILTER DESIGN

Previous work has chosen measurement components such as sensors, illuminants, or filters from discrete sets of existing products. While every distinct measurement provides some information for spectral reconstruction, there is no way to tailor the measurement devices to produce spectra that would prove most useful or target missing information. For example, camera response spectra tend to be very similar, and if one did not have the benefit of laboratory calibration it would be difficult to distinguish metameric spectra based on subtle differences in pixel intensities.

We address this problem by introducing tunable, broadband spectral filters capable of creating a continuous range of transmission spectra that can be calculated analytically. These filters consist of waveplates, or wave retarders, sandwiched between two linear polarizers (Fig. 2). The continuous rotations of these optical elements generate a continuous space of spectra, which grows arbitrarily large as waveplates are added to the system. Conceptually similar to how dispersion and diffraction transform spectral information into spatial information, our filters transform spectral information into polarized information (and back to spectral). We describe the forward model that accurately predicts this optical transformation (Sec. 4.1) and analyze the theoretical spectral gamut that is achievable with different numbers of optical elements and orientation measurements (Sec. 4.2). Finally, we use a simple, greedy algorithm to prescribe a set of filter

TABLE 1: We categorize related linear-reconstruction-based approaches according to whether the illuminant $e(\lambda)$, the filter spectrum $t(\lambda)$, or camera response $c(\lambda)$ are known (●), unknown (○), or not used (–), and whether these are modulated in a continuous (○) or discrete (⊙) manner.

	$r(\lambda)$	$e(\lambda)$	$t(\lambda)$	$c(\lambda)$
Park et al. [8]	○	⊙	–	●
Han et al. [6]	○	⊙	–	●
Hidaka et al. [7]	○	⊙	–	●
Oh et al. [10]	○	○	–	⊙
Chi and Ben-Ezra [9]	○	●	⊙	○
Ours	○	○	⊙	○




Fig. 2: A beam of light is portrayed as a “bundle” of fibers, each with its own wavelength (color), amplitude (diameter), and polarization state (ellipse/line in the cutaway). Incoming light (a) passes through the first linear polarizer (b), which transmits only the polarization component of each fiber aligned with its transmission axis (c). Then, the light passes through the waveplate (d), which causes a *wavelength- and orientation-dependent phase shift* – linear polarization is transformed into another polarization state, generally elliptical polarization (e). In this illustration, long wavelengths are affected least strongly and short wavelengths affected most strongly. Lastly, the light passes through the second linear polarizer (f), or “analyzer”, which again transmits only the polarization component aligned with its transmission axis. The compound effect of the optical elements together is wavelength-dependent transmission (g).

configurations that can achieve good spectral reconstructions given a fixed measurement budget (Sec. 5.1).

4.1 Polarized light

A light wave may be linearly polarized, circularly polarized, or elliptically polarized (the general case), corresponding to the shape that the tip of the light wave’s electric field vector traces as the wave propagates through space. The polarization state of a beam of light comprised of many individual waves can be described by a four-component Stokes vector [42]:

$$\mathbf{s} = \begin{bmatrix} s_0 \\ s_1 \\ s_2 \\ s_3 \end{bmatrix} = \begin{bmatrix} \odot \\ \leftrightarrow - \updownarrow \\ \nearrow - \nwarrow \\ \circ - \circ \end{bmatrix}, \quad (5)$$

where s_0 is the total radiance, s_1 is horizontal (0°) minus vertical (90°) linearly polarized radiance, s_2 is 45° minus 135° linearly polarized radiance, and s_3 is right-hand minus left-hand circularly polarized radiance. Each of s_1 , s_2 , and s_3 lies in the domain $[-s_0, s_0]$. The degree of polarization (DOP) of the light can be quantified as $\sqrt{s_1^2 + s_2^2 + s_3^2} / s_0$, where $\text{DOP} = 0$ is completely unpolarized light, $\text{DOP} = 1$ is completely polarized light, and $0 < \text{DOP} < 1$ is partially polarized light. For example, a beam of light with radiance = 1 that is completely unpolarized would have Stokes vector $\mathbf{s} = [1, 0, 0, 0]^\top$, whereas light that is completely vertically polarized would have Stokes vector $\mathbf{s} = [1, -1, 0, 0]^\top$.

A Stokes vector can be transformed into another by a 4×4 Mueller matrix, which represents an interaction of light with a real-world material, such as a linear polarizer, waveplate, or surface. Our filters require the Mueller matrix of a linear polarizer at angle θ :

$$\mathbf{T}_{\text{LP}}(\theta) = 0.5 \begin{bmatrix} 1 & c & s & 0 \\ c & c^2 & cs & 0 \\ s & cs & s^2 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad (6)$$

and the Mueller matrix of a waveplate with wavelength-dependent phase shift $\Gamma(\lambda)$ at angle θ :

$$\mathbf{T}_W(\Gamma(\lambda), \theta) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & c^2 + s^2 \cos \Gamma & cs(1 - \cos \Gamma) & s \sin \Gamma \\ 0 & cs(1 - \cos \Gamma) & c^2 \cos \Gamma + s^2 & -c \sin \Gamma \\ 0 & -s \sin \Gamma & c \sin \Gamma & \cos \Gamma \end{bmatrix}, \quad (7)$$

where $c = \cos 2\theta$ and $s = \sin 2\theta$. The phase shift $\Gamma(\lambda) = 2\pi \Delta n d / \lambda$ depends on the birefringence Δn of the material, which is the difference in index of refraction between the ordinary and extraordinary rays, thickness d of the waveplate, and wavelength λ . Birefringence causes the orthogonally polarized components of light to travel at different speeds through the material, which creates a phase delay that alters the polarization state of the light.

We can chain together the Mueller matrices of every optical element in the system to realize the full optical effect of the filter:

$$\mathbf{T}_f(\lambda; \Theta_n^j) = \mathbf{T}_{\text{LP}}(\theta_n) \mathbf{T}_W(\Gamma_{n-1}(\lambda), \theta_{n-1}) \cdots \mathbf{T}_W(\Gamma_1(\lambda), \theta_1) \mathbf{T}_{\text{LP}}(0^\circ) \quad (8)$$

where there are $n - 1$ waveplates and the thicknesses d_1, \dots, d_{n-1} and angles $\theta_1, \dots, \theta_{n-1}$ are tunable parameters. The final tunable parameter is the rotation of the analyzer θ_n .² We call a particular choice j of these parameters Θ_n^j a *configuration* of our filter design where n is the number of degrees of freedom.

While the rotation angles are continuous variables, we choose the thicknesses from a discrete set of existing waveplates. Once the parameters are chosen, the transmission spectrum $t(\lambda; \Theta_n^j)$ of a filter can be calculated by passing in radiance-normalized unpolarized light, $\mathbf{s}_1 := [1, 0, 0, 0]^\top$, and isolating the first Stoke’s element:

$$t(\lambda; \Theta_n^j) = \mathbf{s}_1^\top \mathbf{T}_f(\lambda; \Theta_n^j) \mathbf{s}_1. \quad (9)$$

2. Since only the relative rotation of the two polarizers change the filter’s spectrum, we fix one polarizer at 0° and together the two polarizers only add one degree of freedom.

Note that $t(\lambda; \Theta_n^j) > 0$ except in two degenerate cases: either the waveplate produces a phase shift for λ that results in *exactly* an orthogonal polarization state to the analyzer, or the polarizers are orthogonal to each other and the waveplate is parallel to either of the polarizers. Other than those two cases, $t(\lambda; \Theta_n^j)$ is non-zero and provides a value that could potentially aid in reconstruction.

This system assumes that light entering the filter is unpolarized, otherwise the transmission calculation may be inaccurate. We did not run into any related issues in practice; however, special care may need to be taken around illuminants or surfaces that can create polarization, such as LEDs or specular reflections.

Now that we have described the analytic forward model of our filters, we can move on to exploring the space of transmission spectra they are able to achieve. Fig. 1 shows example transmission spectra.

4.2 Gamut of achievable spectra

Let us first formalize the space of achievable spectra F_n . Once a set of waveplates is chosen, there are n independent continuous degrees of freedom: the rotation angles of the waveplates $\theta_1, \dots, \theta_{n-1}$ and the rotation angle of the analyzer θ_n . F_n is the space of transmission spectra $t(\lambda; \Theta_n^j)$ induced by any possible configuration Θ_n^j .

Note that F_n is *infinite, but compact* (in the set theoretical sense). This is easy to see since each $\theta_i \in \Theta_n^j$ is the rotation angle of a filter element and thus bounded in $[0, 2\pi]$, and the Mueller calculus operations that produce a transmission spectrum are non-singular. This lets us define the *spectral gamut* of F_n as a subset of the infinite dimensional space of all possible transmission spectra, analogous to the *color gamut* of a particular display, for example.

To get a sense of the number of linearly independent spectra our filter design can produce, we investigate the dimensionality of F_n . Note that this is not simply equal to n , the number of degrees of freedom of our filter, because the relation (9) between configuration and transmission spectrum is highly non-linear. Unfortunately, it is challenging to analytically find an orthonormal basis for F_n . Instead, we run principle component analysis (PCA) on a discrete set of sample spectra drawn from F_n for $n \in 1, \dots, 100$. We define the effective dimensionality k_n of F_n as the number of linear basis functions derived from PCA that explain 99.9%




Fig. 3: As we increase the number of waveplates, the effective dimensionality of the space of achievable spectra increases. This makes it possible to predict how many useful distinct measurements a particular filter setup can provide.

of the variance in the sampled dataset. The effective basis set is then the first k_n basis functions retrieved via PCA, $f_1(\lambda), \dots, f_{k_n}(\lambda)$. Our empirical experiments confirm that k_n grows arbitrarily large as we add waveplates to our setup (see Fig. 3).

Additionally, the basis functions derived from PCA provide a convenient *distance* metric $\|t(\lambda; \Theta_n^1) - t(\lambda; \Theta_n^2)\|$ between two elements in F_n . The Euclidean distance in this low-dimensional space fundamentally measures how “different” the transmission spectra produced by two distinct configurations Θ_n^1 and Θ_n^2 are. By plotting the achievable transmission spectra F_n in our basis space, giving us a gamut plot as shown in Fig. 4, we can verify that spectra with a large distance between them are, in fact, substantially different.

We have shown that the gamut of achievable transmission is theoretically wide enough to serve as a good basis to measure a wide range of unknown spectra. Still, it is not obvious how to pick a discrete set of filters from F_n to most accurately reconstruct an arbitrary measurement spectrum. We discuss our approach to solve this problem next.

5 METHODS

5.1 Filter choice

One advantage of a system with a small discrete set of choices is that it is straightforward to determine the optimal subset to use. However, our system has both discrete (which waveplates to use and in what order) and continuous (rotation angle) variables. If we were to run an optimization to find the

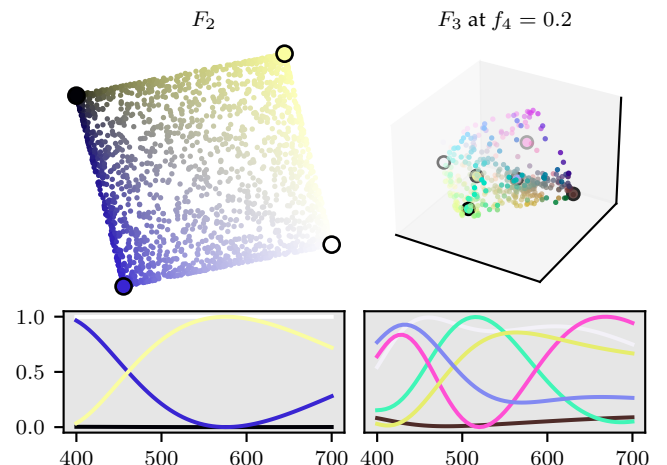


Fig. 4: In the top row we plot points sampled in F_2 (left) and from a slice through F_3 (right). Recall that each point in F_n is a transmission spectrum $t(\lambda; \Theta_n^j)$. To visualize these spectra we plot their projection to sRGB space. F_2 has the effective dimensionality $k_2 = 2$ and hence is easy to visualize. For F_3 , $k_3 = 4$ and we have 4 basis functions: f_1, f_2, f_3, f_4 . For visualization purposes, we choose to vary the f_1, f_2 and f_3 components along the x, y and z-axes respectively and keep the remaining component fixed. In the bottom row we show the full spectra for the extreme points highlighted in the respective spaces above. Note that the spectra are different and we can produce a larger number of distinct spectra in F_3 than in F_2 .

optimal parameters that minimize the reconstruction error of a large reflectance dataset, this would be a mixed integer non-linear optimization problem, which are notoriously difficult to solve and for which few solvers are readily available. Alternately, we could chop up the continuous parameters into discrete sets and try to run a global optimization such as an evolutionary algorithm to find the optimal subset of filters; however, this would be intractable time-wise since the dataset of achievable filters is too large. For example, [9] use an evolutionary algorithm to choose an optimal subset of 8 to 16 theatrical filters out of 230 total, and their optimizations ran for a few days. For our problem this approach is not practical since even a coarse five-degree increment discretization of a two-waveplate system would create a dataset of 186,624 filters.

Instead, we draw from our analysis in Sec. 4.2 to devise a simple strategy that we found is both efficient and produces good results. Given a fixed measurement budget, we would like to choose filters so as to fill in as much missing information in the unknown spectrum as possible. We do this by choosing filters whose spectra are most distant or stratified within the space of all possible filters – a measure we can quantify using the low-dimensional basis space described in Sec. 4.2. This will naturally provide measurements that are distinct from one another in different areas of the visible spectrum, providing distinguishing information to hone our estimate of the unknown spectrum.

We first create a large random dataset of filter candidates (typically 5,000) and project them down to a low-dimensional space using principal component analysis. When building the dataset, we reject any randomly built filter with an average transmission below 20%. Given a fixed budget of M measurements, we run a fast, approximate, greedy algorithm [43], [44] to find a cluster of M points in this low-dimensional filter space that maximize their minimum intracluster distance in Euclidean space. We first choose any filter in the dataset, then run a search for the filter in the dataset with the maximum distance from the initial point and add it to our cluster. We repeat this process iteratively, at each stage finding the filter in the dataset with the maximum distance from *all* previously chosen filters, until the fixed cluster size has been reached (see Alg. 1). The combined runtime of both creating a random dataset of 5,000 filters and running the greedy algorithm to choose a subset of two to ten filters is under ten seconds.

To validate our stratified choice algorithm, we compare the performance of its choice of filter set against 5,000 randomly chosen filter sets on the reconstruction loss of a set of 200 noise-perturbed measurements ($\sigma = 0.05$ with $0 < p < 1$) of Munsell chip reflectance spectra for up to 20 measurements (Fig. 5). We also compare to the strategy of choosing filters whose transmission spectra produce the lowest condition number when stacked in a matrix, which is analogous to the choice strategies from prior work [9], [10]. Our stratified choice strategy always achieves error within one standard deviation of the random sets, and in most cases achieves lower error than both the lowest condition number strategy and the lower bound of the random sets' standard deviation. Even though there are a few cases where the condition number strategy outperforms ours, its behavior is highly erratic and it often yields reconstruction loss above




Fig. 5: Reconstruction loss of a random choice strategy, lowest condition number strategy, and our strategy for choosing filter sets given a fixed measurement budget. For random choice, the mean reconstruction loss over 5,000 trials is plotted with error bars for standard deviation. Dashed best-fit lines are shown for the lowest condition number strategy and our strategy. Our strategy consistently achieves error within one standard deviation of the random set, and in most cases performs consistently better than the other two strategies.

one standard deviation of the random sets. In contrast, our strategy shows more pronounced improvement over both other strategies as the number of measurements increases. This becomes especially relevant when we take advantage of spatial multiplexing of our filters, which allows us to take many measurements in a single photograph. Taking several photographs with a fixed array of filters expands the number of measurements rapidly.

To further explore the relationship between sample stratification and reconstruction loss, in Fig. 6 we calculate the minimum intracluster distance for each of the 5,000 randomly chosen filter sets (binned into box plots), the lowest condition number strategy (in green), and our algorithm's choice (in red) and plot their reconstruction loss for fixed numbers of measurements (M). Although the median reconstruction loss of the random filters remains relatively steady as the minimum intracluster distance increases, there is a clear inverse relationship between the variance of reconstruction loss and minimum intracluster distance. This suggests that while our stratified choice strategy may not always achieve the absolute lowest reconstruction loss, it is likely to be a

ALGORITHM 1: Filter Choice

```

Input :  $S$ , a set of filter spectra
          $M$ , a fixed number of measurements
Output:  $S'$ , a subset of  $S$  of size  $M$ 
 $S' \leftarrow \{S[0]\}$  // initialize  $S'$  to any element of  $S$ 
for  $i = 2, \dots, M$  do
    find  $p \in S \setminus S'$  which maximizes  $\min_{x \in S'} \|x - p\|$ 
    // find the point  $p$  which is most
    // distant from all of the already chosen
    // filters
     $S' \leftarrow S' \cup \{p\}$ 
end
return  $S'$ 

```

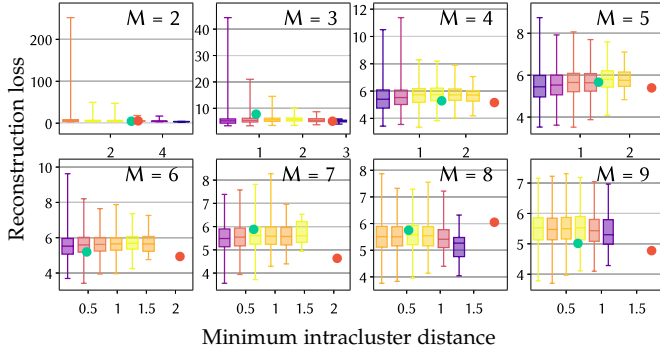


Fig. 6: Reconstruction loss of 5,000 randomly chosen filter sets (binned into boxplots) versus the filter set chosen by our stratified filter choice strategy (in red) and a minimum condition number strategy (in green) for several fixed measurement budgets (M). Each filter set has been projected down to four-dimensional space and the minimum distance between any two points within the set is plotted on the x-axis, which have been binned for easy comparison. Each box plot displays the minimum, first quartile, median, third quartile, and maximum value for the data points in that bin. The median values of each box plot are mapped to a color scale from purple (low) to yellow (high).

reliably good choice.

We used two-waveplate filters in all our experiments, since we determined through experimentation that the four-dimensional space of achievable spectra with two-waveplate filters is sufficient for spectra recovery.

5.2 Solving for unknown spectra

In order to find the spectrum $u(\lambda)$ which minimizes the loss in Eq. (4), we transform each continuous spectrum into a vector by sampling it at a discrete set of wavelengths. We denote the number of elements in the resulting spectral vector λ .

We could in theory solve for unknown spectrum $\mathbf{u} \in \mathbb{R}^\lambda$ by finding the least-squares solution to the following linear system:

$$\begin{bmatrix} \text{---} \mathbf{m}_1 \text{---} \\ \text{---} \mathbf{m}_2 \text{---} \\ \vdots \\ \text{---} \mathbf{m}_M \text{---} \end{bmatrix} \mathbf{u} = \begin{bmatrix} p_1 \\ p_2 \\ \vdots \\ p_M \end{bmatrix} \quad (10)$$

$\underbrace{\hspace{10em}}_{M \times \lambda} \quad \underbrace{\hspace{2em}}_{\lambda \times 1} \quad \underbrace{\hspace{5em}}_{M \times 1}$

where \mathbf{m} is any spectral measurement vector, e.g. $\mathbf{m}_1 = \mathbf{c}_R \odot \mathbf{t}_1 \odot \mathbf{e}$ and p_1 is the pixel's red intensity value under a filter \mathbf{t}_1 and illuminant \mathbf{e} (\odot is component-wise multiplication). We denote the spectral measurement matrix on the left-hand side of Eq. (10) as $\mathbf{M} \in \mathbb{R}^{M \times \lambda}$ and the pixel vector on the right-hand side as \mathbf{p} . Unless we have a very large number of measurements, this system will be underdetermined given that $\lambda \gg M$ (or arbitrarily determined by the size of λ).

It is well established that real-world reflectance spectra live in a low-dimensional space [45]. We can therefore transform Eq. (10) into an overdetermined system by first

representing \mathbf{u} using a set of basis vectors \mathbf{b}_i computed from a large database of 1,250 Munsell chips using PCA:

$$\mathbf{u} \approx \mathbf{B}\mathbf{a} = \begin{bmatrix} | & | & & | \\ \mathbf{b}_1 & \mathbf{b}_2 & \cdots & \mathbf{b}_8 \\ | & | & & | \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_8 \end{bmatrix}. \quad (11)$$

Like other related works [6]–[8], [10], we choose to use the first eight basis functions, which account for over 99.7% of the total variance. Inserting this into Eq. (10), we obtain the modified system $(\mathbf{M}\mathbf{B})\mathbf{a} = \mathbf{p}$,

$$\begin{bmatrix} \text{---} \mathbf{m}_1 \text{---} \\ \text{---} \mathbf{m}_2 \text{---} \\ \vdots \\ \text{---} \mathbf{m}_M \text{---} \end{bmatrix} \begin{bmatrix} | & | & & | \\ \mathbf{b}_1 & \mathbf{b}_2 & \cdots & \mathbf{b}_8 \\ | & | & & | \end{bmatrix} \mathbf{a} = \begin{bmatrix} p_1 \\ p_2 \\ \vdots \\ p_M \end{bmatrix}, \quad (12)$$

$\underbrace{\hspace{15em}}_{M \times 8}$

which we can now easily overconstrain with just a few measurements and solve for the eight unknowns in a using least squares.

We also impose a smoothness constraint, since reflectance spectra in the real world transition smoothly between wavelengths, and restrict the unknown spectra to be between zero and one (see Oh et al. [10] for details on the matrix formulation of these constraints).

5.3 Spectral resolution

The choice of λ , the number of elements which represent a spectral vector, will determine the spectral resolution of the system. A natural choice is to use the spectral resolution of the basis functions used to construct the spectra. The basis functions for reflectances, illuminants, sensor responses, etc. are built from real-world datasets reported in varying spectral resolutions that must be converted to a common resolution before they can be combined in a linear system. We chose the lowest resolution of these sources as the common resolution for our reconstructions, which was 10 nm intervals in the range 400 nm to 700 nm.

5.4 White balance correction

We initially borrowed Oh et al. [10]'s method to estimate unknown surface reflectances and an unknown illuminant simultaneously by alternating between solving for the reflectances and solving for the illuminant (which follows the same logic, just switching which quantity is known versus unknown). However, we noticed in our initial experiments that some of the sRGB projections of the acquired spectra did not match the sRGB projections of the ground truth, indicating that our system may be biased. For example, we assume that the polarizing and clear filters have wavelength-independent transmission, which may not hold true in reality. To account for this bias, we developed a preprocess for white balance correction using the known reflectance spectra of all the ColorChecker squares. During measurement, we take a white balance photo of the ColorChecker *with the filter optical elements aligned at zero degrees*, which in theory should allow 50% wavelength-independent transmission (the first linear polarizer attenuates incoming unpolarized light by

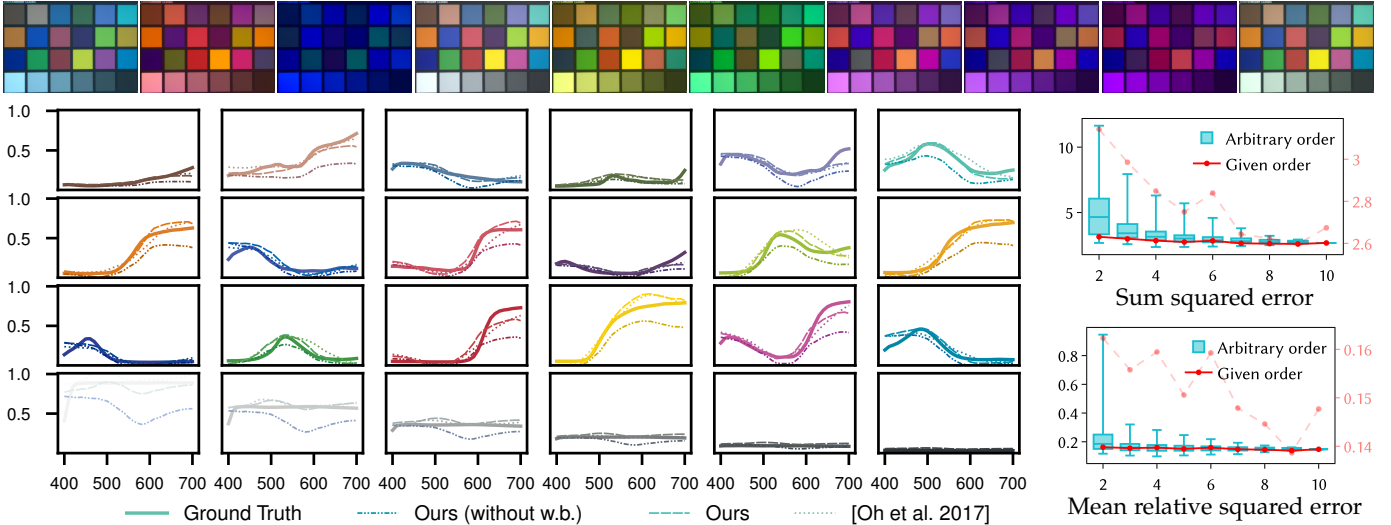


Fig. 7: Recovered Calibrite ColorChecker reflectance spectra without and with white balance correction show good agreement with ground truth (the color of each line is an sRGB projection of the recovered spectrum). The top row shows the measurement photos used in reconstruction. We show the error for our recovered spectra (with white balance) at a fixed number of measurements in the filter order chosen by our algorithm (in red, with an expanded scale in lighter red) and also compare to the errors achieved by all possible combinations of the measurement photos (box plots). No matter which photos are used, the variance of the reconstruction loss decreases as more measurements are added.

50%). Then, instead of estimating the unknown reflectance and unknown illuminant simultaneously, we estimate a white balance correction factor for the unknown illuminant and sensor responses simultaneously by alternating least squares (using our camera model’s reported response curves as an initial guess), treating the reflectance spectra of the ColorChecker squares as known. Any wavelength-dependent bias should be baked into the recovered illuminant and sensor responses. From here, we treat the illuminant and sensor responses as knowns and estimate the unknown scene spectra. We show in Sec. 6 how this correction improves the results.

6 SPECTRAL REFLECTANCE RECOVERY

6.1 Physical capture setup

We took advantage of the wide array of photography filters designed to screw into the lenses of ordinary digital cameras in order to build our prototype. We took all measurements with a Nikon D5100 DSLR camera, for which the sensor responses are reported online [46]. For linear polarizers, we used 52 mm Hoya linear polarizing filters, which are designed to rotate freely. We attached each of our waveplates (described below) to a 52 mm Nikon clear filter. These clear filters can be “rotated” by being unscrewed partway. We labeled angles of rotation relative to the linear polarizing axis or the fast axis in five-degree increments all around the outer rings of each filter (see Fig. 1). These filters have threads on both sides, allowing them to be stacked on top of one another to create our spectral filters.

We choose to use clear packing tapes as cheap waveplates, whose birefringent properties arise from long polymers stretched along the length of the tape. We used two types of tape in our experiments, Scotch Sure Start packing tape ($d = 66.04 \mu\text{m}$) and Scotch Heavy Duty packing tape

($d = 78.74 \mu\text{m}$), thicknesses reported by the manufacturer. To calculate the birefringence of our tapes, we used the method of Beléndez et al. [47]: we take a set of measurements of a tape sandwiched between two linear polarizers while illuminated by laser pointers of known wavelengths. We provide the details of this birefringence estimation setup in our supplemental. Using this calibration method, we retrieved an estimated birefringence of 0.00872 for the Sure Start tape and 0.00936 for the Heavy Duty tape.

We provide detailed instructions for reproducing our capture setups in the supplemental material.

6.2 Recovery of ColorChecker Spectra

We validate our system by recovering the reflectance spectra of the squares of a Calibrite ColorChecker with known reflectance spectra [48] using measurements of the ColorChecker taken through 10 distinct filters chosen by our filter choice algorithm. Our spectral reconstructions with and without the white balance correction step are shown in Fig. 7.

Throughout, we define sum squared error (SSE) as $\sum_{i=1}^s \sum_{j=1}^{\lambda} (u_{i,j} - \hat{u}_{i,j})^2$, where s is the number of ColorChecker squares, u_i is the recovered value, and \hat{u}_i is ground truth. To equalize the error contribution of all the ColorChecker squares, we also calculate the mean relative squared error (MRSE), defined as $= \frac{1}{s} \frac{1}{\lambda} \sum_{i=1}^s \sum_{j=1}^{\lambda} (u_{i,j} - \hat{u}_{i,j})^2 / (\hat{u}_{i,j}^2 + \epsilon)$, where $\epsilon = 1 \times 10^{-8}$ avoids division by zero.

Our system is able to recover the reflectance spectra of most of the squares faithfully, but struggles the most with the lightest gray squares. The shape of our reconstructions and their deviations from ground truth closely match those of best-case-scenario reconstructions using our method in simulation (provided in our supplemental material) and also those of prior linear-reconstruction-based work [10], which suggests that the error is mainly due to the inherent limitations of the linear-reconstruction-based problem formulation.

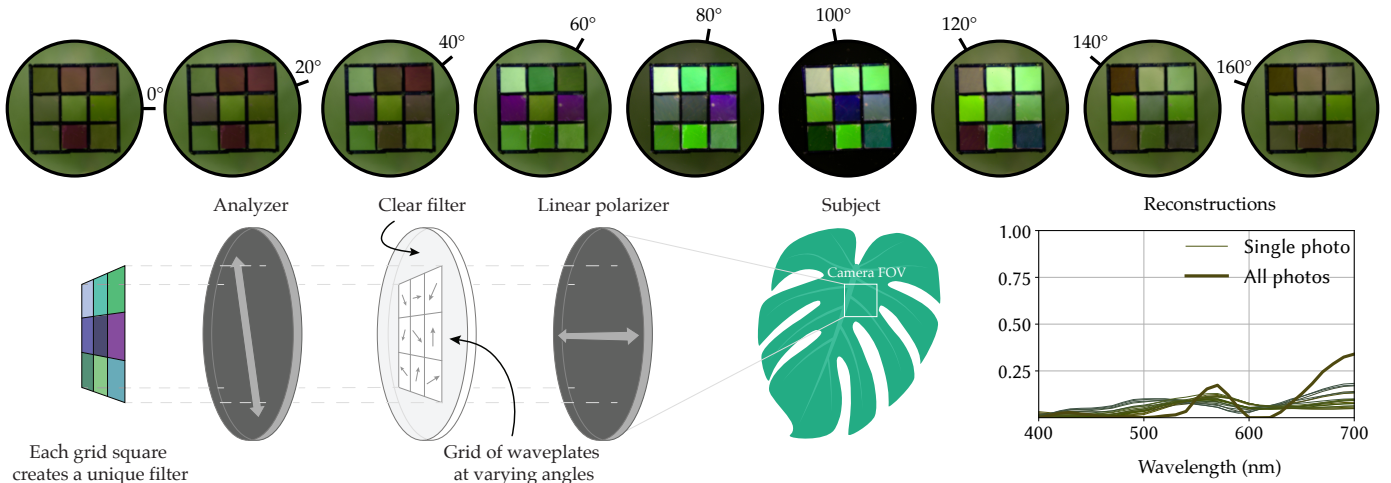


Fig. 8: The rotation of the analyzer around a 3×3 grid of waveplates fixed at various orientations produces a “kaleidoscope”-like effect (*top row*). A portion of the leaf of interest occupies the camera’s entire field of view. Each measurement photo could be used individually as a single-shot strategy toward reconstructing the spectral reflectance of the *Monstera deliciosa* leaf (*bottom-right, thin lines*), or all together (*bottom-right, thick line*).

We compare our results to Oh et al. [10], who propose a do-it-yourself hyperspectral imaging setup using three different digital cameras. Their approach is highly related to our own, utilizing the same mathematical framework with the only difference being how we each create distinct measurement spectra, $m(\lambda)$. The spectra graphs in Fig. 7 show qualitatively that our reconstructions are comparable, slightly better in some places and slightly worse in others, than those of Oh et al. [10]. We achieve similar SSE (ours: 2.67, theirs: 2.84) and MRSE (ours: 0.148, theirs: 0.110). Our system is less expensive (requiring just one digital camera and the filter elements, compared to three digital cameras), and also does not require calibrated camera response curves or calculating image homographies for registration. On the other hand, Oh et al. [10]’s system requires less setup by the user.

We also look at how the error of reconstruction behaves as we vary the number of measurement photos used (from 2 to 10) and which measurement photos are used. In Fig. 7 (right), for a fixed measurement budget $M \in [2, 10]$, we calculate the reconstruction error of up to 500 randomly chosen sets of size M of our measurement photos (“Arbitrary order”). The top set of boxplots shows the distribution of sum squared error (SSE) for this dataset and the bottom set shows the mean relative squared error (MRSE). We compare these random permutations to the order given by our filter choice algorithm (“Given order”) – due to the greedy nature of our algorithm, each additional measurement adds one new filter configuration to the already existing set. The variance of the error amongst all possible permutations is high initially and tails off steadily as measurements are added to the system, converging to a steady error value by 10 measurements. The SSE and MRSE achieved by our choice algorithm lie consistently below the median errors, although it does not always achieve the absolute lowest error. This suggests that the order provided by our filter choice algorithm is a reliably good one, and that adding more measurements should reliably reduce the overall error.

One possible factor that may account for some combinations of measurements performing better than others is that some measurements may be more affected by user error, e.g. the angle of rotation may be off.

Our supplemental material provides more analysis of possible sources of error on reconstructions. We found that our reconstructions are most sensitive to errors in the estimate of tape birefringence and thickness, and relatively insensitive to errors in angles of rotation. Therefore, it is especially important to mitigate error when performing the birefringence calibration step. With carefully labeled angles, estimating the angles of rotation should not be a significant source of error in the reconstruction.

6.3 Recovery of an Unknown Spectrum by Mosaic Filter

One drawback of attaching a single tape to each clear filter is that it takes time to set every element of the filter stack to the proper angle before every measurement. We can trade off temporal multiplexing for spatial multiplexing by creating a grid of differently oriented waveplates on a single “mosaic” filter. This technique allows us to create a *single-shot* approach, where many measurements can be taken with one photograph at the cost of the unknown spectrum of interest needing to span the field of view. We demonstrate using a mosaic filter with nine waveplate configurations to acquire the reflectance spectrum of a *Monstera deliciosa* leaf in Fig. 8. We show the resulting set of nine measurements for a range of analyzer angles in the top row, illustrating the kaleidoscope-like morphing of colors as the analyzer is rotated. The reconstructed spectra that result from using each measurement photo individually as a single-shot strategy are shown in Fig. 8 (bottom-right, thin lines), as well as a multi-shot reconstruction that uses all measurement photos together, rotating the analyzer from 0° to 180° in increments of 10° (bottom-right, thick line). The recovered reflectance spectra clearly mirror the known absorption spectra of chlorophyll *a* and chlorophyll *b*, the dominant pigments in green plant leaves, whose absorption is minimal between

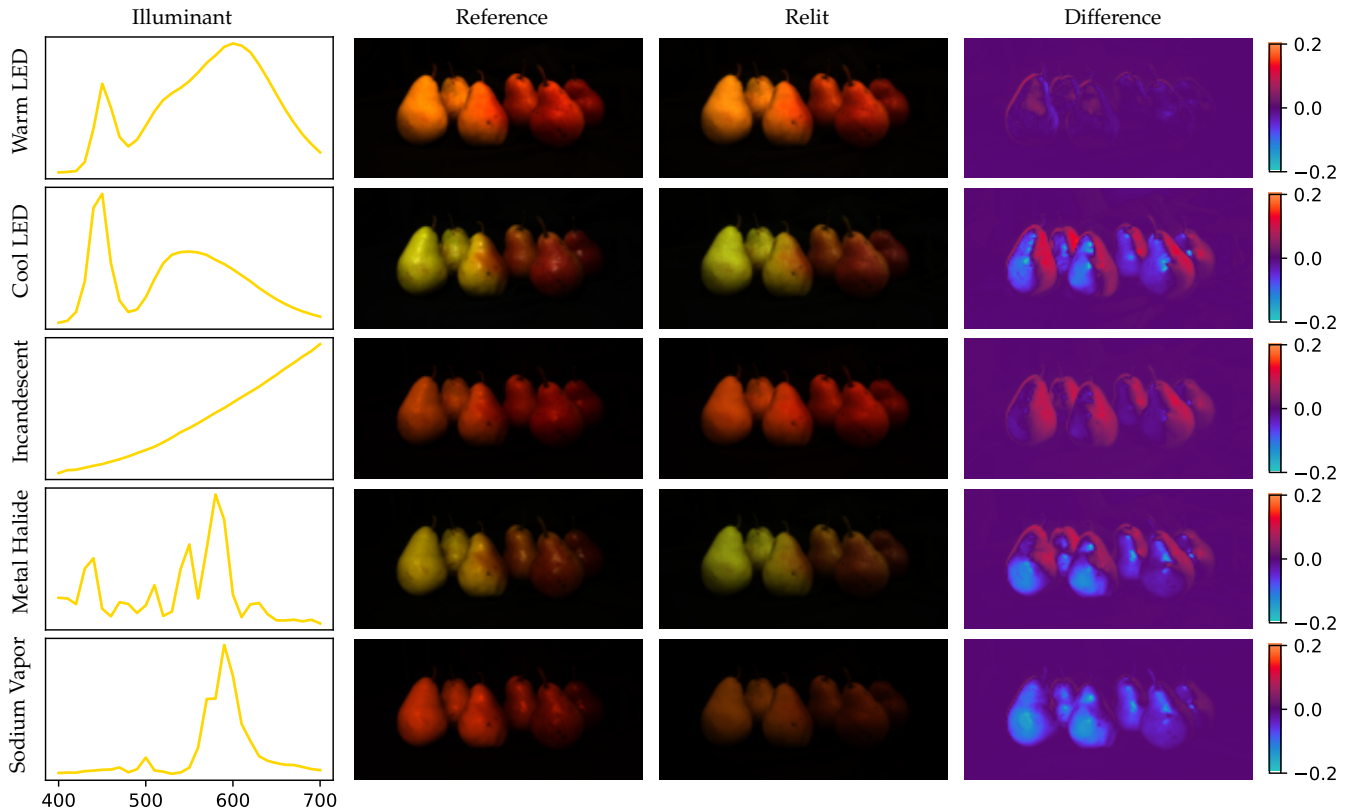


Fig. 9: Comparison of relighting a scene physically versus virtually. A variety of illuminants (emission spectra in left column) are used to relight a scene in the real world (Reference, middle left) and their spectra are multiplied by our reconstructed reflectance spectra and projected down to RGB using our estimated sensor responses (Relit, middle right). The difference per pixel is shown at right. The top row relights the scene using its original illuminant and shows very good agreement with the reference. Our relit results are able to recreate the color palette of the reference results – the differences in highlight and shadow shapes arise due to the physical properties of the light sources.

500 and 600 nm [49]. The single-shot reconstructions are smoother than the multi-shot reconstruction, demonstrating that as more measurements are added to the system, the relative importance of the smoothing term in the linear system decreases. We leave the development of an adaptive smoothing term that compensates for this effect to future work.

One could extend this method to efficiently take even more measurements by, for example, taking a video of the continuous rotation of the analyzer and extracting the frames, or shrinking the grid size to fit more waveplate squares on the clear filter. Following that thought to its extreme, one could imagine creating a Bayer-like mosaic where the grid is on the scale of the camera sensor pixels. This would overcome the limitation of needing the subject to span a large portion of the image.

6.4 Recovery of Unknown Scene Spectra and Relighting

One application of hyperspectral imaging is relighting: separating a scene’s spectra from its illuminant in order to visualize it under a different illuminant. Relighting also provides a layer of validation for our method in that we would expect to be able to “relight” our retrieved spectra virtually under the given spectrum of any arbitrary illuminant and have it match a captured image of the same scene under that illuminant.

For this relighting experiment, we first took 10 measurements of a scene using the filter configurations provided by our filter choice algorithm under warm LED lights. While we could use our white balance method described in Sec. 5.4 to estimate the sensor responses and illuminant simultaneously, there is no reliable way to control where the spectral correction factors are distributed among the estimated spectral responses or illuminant – this does not matter for the purposes of retrieving the reflectance spectra, but in this case we need to distribute any such corrections into the sensor responses and leave the illuminant spectrum unaltered. Therefore, we acquire a ground truth illuminant spectrum with a Sekonic C-7000 SpectroMaster spectrometer and estimate sensor responses, along with any necessary correction factor, using this ground truth illuminant. We downsized each measurement photo to 205×309 and solved for the reflectance spectra at each pixel, treating the ground truth illuminant and estimated sensor responses as known. We then relit the scene both in simulation and in real life to assess the accuracy of our reconstructions.

The results of our relighting experiment are shown in Fig. 9. We tried to use a variety of illuminants in order to test the robustness of our reconstructions – their emission spectra acquired with our spectrometer are shown in the leftmost column, normalized to their maximum value. For reference images, we took an image of our scene without any filters

over the camera lens under each illuminant. We tried as much as possible to place the light sources in the same position relative to our scene, although the original measurement photos were taken under more spatially distributed overhead lighting and create less intense highlights and shadows. For our virtually relit images, we projected the product of the illuminant spectrum and our reconstructed reflectance spectra to RGB using our estimated sensor responses. The absolute difference in pixel value between the reference and relit images (averaged among color channels) is shown in the rightmost column. We adjusted the relative intensities of the ground truth illuminant spectra to match the relative exposure times of the reference images when performing the relighting. The top row relights the scene using the original illuminant used for reconstruction, and gives an idea of the minimum possible error our method could achieve – the agreement is very good, with only slight error at silhouette edges and highlights. The rest of the images clearly show slight differences in the locations and intensities of highlights and shadows, and much of this difference can be attributed to the physical shapes, locations, and intensities of the light sources varying between photos. Generally, the colors and their relative intensities match well between the reference and relit images.

7 DISCUSSION, LIMITATIONS, & FUTURE WORK

We presented an approach for do-it-yourself hyperspectral imaging utilizing a novel design for broadband spectral filters that transform spectral information to and from polarized information. These spectral filters, comprised of polarizers and waveplates, can be tuned to produce a continuous range of transmission spectra simply by rotating their optical elements. We showed that the dimensionality of the space of achievable transmission spectra grows arbitrarily large as waveplates are added to the filter, and we addressed the practical aspects of choosing a set of filters from this continuous space to use for measurements.

While our choice of filters from a continuum is reasonable and fast, it does not always achieve the global minimum reconstruction loss when compared to a large random set. Exploring other strategies for filter choice that perform better in a reasonable time frame would be useful not only for this particular system, but other linear-reconstruction-based methods that choose measurement spectra from discrete sets.

We also noticed some chromatic aberrations toward the outer edges of our filters that suggest the transmission spectrum may not be homogeneous throughout. We suspect this is due to light toward the edges of the filter entering at an angle, violating our assumption of light entering normal to the filter plane, which would cause it to doubly refract. Further investigation via realistic rendering of birefringent materials may yield insights into why this occurs.

Lastly, while we demonstrated the validity of our filter design on a low-cost prototype, our design would be equally valid using polarizers and waveplates of any quality. The mechanism of our filters is highly related to how LCDs work – each pixel consists of a liquid crystal sandwiched between two polarizers whose polarizing properties are altered with rapid electrical signals. Using liquid crystals with voltage-modulated properties would turn the discrete choice of

waveplate properties into a continuous parameter that would further expand the space of achievable spectra. Adapting such a setup for our use case could allow for acquisition speeds well beyond mechanical limits and also allow pixels to operate independently of one another, perhaps coordinating to gather different sets of information in parallel (see [40] for an example of this in concurrent work). Our framework opens up many exciting avenues for further exploration that are variations of our core concept with varying hardware implementations.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their suggestions on improving the paper. We would also like to thank Jeremy Salesin for consulting on practical details of building our hardware prototypes. This work was generously supported by NSF award 1844538 and a Neukom Institute CompX faculty grant.

REFERENCES

- [1] M. Nimier-David, D. Vicini, T. Zeltner, and W. Jakob, "Mitsuba 2: A retargetable forward and inverse renderer," *Proc. SIGGRAPH Asia*, vol. 38, no. 6, pp. 203:1–203:17, Nov. 2019.
- [2] L. Fascione, J. Hanika, M. Leone, M. Droske, J. Schwarzhaupt, T. Davidovič, A. Weidlich, and J. Meng, "Manuka: A batch-shading architecture for spectral path tracing in movie production," *ACM TOG*, vol. 37, no. 3, Aug. 2018.
- [3] S.-H. Baek, I. Kim, D. Gutierrez, and M. H. Kim, "Compact single-shot hyperspectral imaging using a prism," *Proc. SIGGRAPH Asia*, vol. 36, no. 6, Nov. 2017.
- [4] R. Habel, M. Kudenov, and M. Wimmer, "Practical spectral photography," *CGF*, vol. 31, 2012.
- [5] Y. Zhao, H. Guo, Z. Ma, X. Cao, T. Yue, and X. Hu, "Hyperspectral imaging with random printed mask," in *Proc. IEEE CVPR*. New York, NY, USA: IEEE, Jun. 2019.
- [6] S. Han, I. Sato, T. Okabe, and Y. Sato, "Fast spectral reflectance recovery using dlp projector," in *Computer Vision – ACCV 2010*. Springer Berlin Heidelberg, 2011.
- [7] H. Hidaka, Y. Monno, and M. Okutomi, "Spectral reflectance estimation using projector with unknown spectral power distribution," 2020. [Online]. Available: <https://arxiv.org/abs/2012.10083>
- [8] J.-I. Park, M.-H. Lee, M. D. Grossberg, and S. K. Nayar, "Multispectral imaging using multiplexed illumination," in *Proc. ICCV*. New York, NY, USA: IEEE, Oct. 2007, pp. 1–8.
- [9] C. Chi and M. Ben-Ezra, "Spectral probing: Multi-spectral imaging by optimized wide band illumination," in *Proceedings of the First International Workshop on Photometric Analysis For Computer Vision – PACV 2007*. INRIA, Oct. 2007.
- [10] S. W. Oh, M. S. Brown, M. Pollefeys, and S. J. Kim, "Do it yourself hyperspectral imaging with everyday digital cameras," in *Proc. IEEE CVPR*. New York, NY, USA: IEEE, 2016.
- [11] E. T. Institute, "Polarized-light mosaic," 2020. [Online]. Available: <https://www.exploratorium.edu/snacks/polarized-light-mosaic>
- [12] Z. Tian, Y.-L. Wei, W.-N. Chang, X. Xiong, C. Zheng, H.-M. Tsai, K. C.-J. Lin, and X. Zhou, "Augmenting indoor inertial tracking with polarized light," in *Proc. Mobile Systems, Applications, and Services*, ser. MobiSys '18. New York, NY, USA: ACM Press, 2018.
- [13] M. E. Gehm, R. John, D. J. Brady, R. M. Willett, and T. J. Schulz, "Single-shot compressive spectral imaging with a dual-disperser architecture," *Optics Express*, vol. 15, no. 21, Oct. 2007.
- [14] A. Wagadarikar, R. John, R. Willett, and D. Brady, "Single disperser design for coded aperture snapshot spectral imaging," *Applied Optics*, vol. 47, no. 10, Apr. 2008.
- [15] G. R. Arce, D. J. Brady, L. Carin, H. Arguello, and D. S. Kittle, "Compressive coded aperture spectral imaging: An introduction," *IEEE Signal Processing Magazine*, vol. 31, no. 1, 2014.
- [16] M. H. Kim, T. A. Harvey, D. S. Kittle, H. Rushmeier, J. Dorsey, R. O. Prum, and D. J. Brady, "3D imaging spectroscopy for measuring hyperspectral patterns on solid objects," *ACM TOG*, vol. 31, no. 4, Jul. 2012.

- [17] H. Du, X. Tong, X. Cao, and S. Lin, "A prism-based system for multispectral video acquisition," in *Proc. ICCV*. New York, NY, USA: IEEE, 2009.
- [18] R. Hostettler, R. Habel, M. Gross, and W. Jarosz, "Dispersion-based color projection using masked prisms," *CGF (Proc. Pacific Graphics)*, vol. 34, no. 7, pp. 329–338, Oct. 2015.
- [19] W. R. Johnson, D. W. Wilson, W. Fink, M. S. H. M.D., and G. H. Bearman, "Snapshot hyperspectral imaging in ophthalmology," *Journal of Biomedical Optics*, vol. 12, no. 1, 2007.
- [20] T. Okamoto, A. Takahashi, and I. Yamaguchi, "Simultaneous acquisition of spectral and spatial intensity distribution," *Applied Spectroscopy*, vol. 47, no. 8, Aug. 1993.
- [21] V. Saragadam and A. C. Sankaranarayanan, "Krism—krylov subspace-based optical computing of hyperspectral images," *ACM TOG*, vol. 38, no. 5, Oct. 2019.
- [22] D. S. Jeon, S.-H. Baek, S. Yi, Q. Fu, X. Dun, W. Heidrich, and M. H. Kim, "Compact snapshot hyperspectral imaging with diffracted rotation," *Proc. SIGGRAPH*, vol. 38, no. 4, Jul. 2019.
- [23] K. Monakhova, K. Yanny, N. Aggarwal, and L. Waller, "Spectral diffusercam: lensless snapshot hyperspectral imaging with a spectral filter array," *Optica*, vol. 7, no. 10, Oct. 2020.
- [24] J. Craven-Jones, M. W. Kudenov, M. G. Stapelbroek, and E. L. Dereniak, "Infrared hyperspectral imaging polarimeter using birefringent prisms," *Applied Optics*, vol. 50, no. 8, Mar. 2011.
- [25] T.-H. Tsai and D. J. Brady, "Coded aperture snapshot spectral polarization imaging," *Applied Optics*, vol. 52, no. 10, Apr. 2013.
- [26] A. Perri, B. E. N. de Faria, D. C. T. Ferreira, F. Preda, D. Polli, A. M. de Paola, D. Comelli, G. Valentini, G. Cerullo, and C. Manzoni, "A hyperspectral camera based on a birefringent ultrastable common-path interferometer," in *Conference on Lasers and Electro-Optics Europe & European Quantum Electronics Conference*. Optical Society of America, 2019.
- [27] C. Bai, J. Li, Y. Xu, H. Yuan, and J. Liu, "Compact birefringent interferometer for fourier transform hyperspectral imaging," *Optics Express*, vol. 26, no. 2, Jan. 2018.
- [28] A. C. Sankaranarayanan, V. Saragadam, V. Rengarajan, R. Tadano, T. Zhuang, H. Oyaizu, and J. Murayama, "Programmable spectral filter arrays for hyperspectral imaging," 2021. [Online]. Available: <https://arxiv.org/abs/2109.14450>
- [29] D. L. MacAdam, "Maximum visual efficiency of colored materials," *Journal of the Optical Society of America*, vol. 25, no. 11, Nov. 1935.
- [30] —, "The theory of the maximum visual efficiency of colored materials," *Journal of the Optical Society of America*, vol. 25, no. 8, Aug. 1935.
- [31] B. Smits, "An RGB-to-spectrum conversion for reflectances," *JGT*, vol. 4, no. 4, 1999.
- [32] H. Otsu, M. Yamamoto, and T. Hachisuka, "Reproducing spectral reflectances from tristimulus colours," *CGF*, vol. 37, no. 6, 2018.
- [33] J. Meng, F. Simon, J. Hanika, and C. Dachsbacher, "Physically meaningful rendering using tristimulus colours," *CGF*, vol. 34, no. 4, 2015.
- [34] W. Jakob and J. Hanika, "A low-dimensional function space for efficient spectral upsampling," *Proc. EG*, vol. 38, no. 2, Mar. 2019.
- [35] I. Choi, D. S. Jeon, G. Nam, D. Gutierrez, and M. H. Kim, "High-quality hyperspectral reconstruction using a spectral prior," *Proc. SIGGRAPH Asia*, vol. 36, no. 6, pp. 218:1–218:13, Nov. 2017.
- [36] L. Wang, C. Sun, Y. Fu, M. H. Kim, and H. Huang, "Hyperspectral image reconstruction using a deep spatial-spectral prior," in *Proc. IEEE CVPR*. New York, NY, USA: IEEE, Jun. 2019.
- [37] M. Paoletti, J. Haut, J. Plaza, and A. Plaza, "Deep learning classifiers for hyperspectral imaging: A review," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 158, 2019.
- [38] Z. Shi, C. Chen, Z. Xiong, D. Liu, and F. Wu, "HSCNN+: Advanced CNN-based hyperspectral recovery from RGB images," in *Proc. IEEE CVPR*. New York, NY, USA: IEEE, 2018.
- [39] J. Zhang, Y. Meuret, X. Wang, and K. A. G. Smet, "Improved and robust spectral reflectance estimation," *LEUKOS*, vol. 0, no. 0, 2020.
- [40] L. Huang, R. Luo, X. Liu, and X. Hao, "Spectral imaging with deep learning," *Light Sci. Appl.*, vol. 11, no. 1, p. 61, Mar. 2022.
- [41] A. Langlands and L. Fascione, "Physlight: An end-to-end pipeline for scene-referred lighting," in *ACM SIGGRAPH Talks*, 2020.
- [42] G. G. Stokes, *On the Composition and Resolution of Streams of Polarized Light from different Sources*, ser. Cambridge Library Collection - Mathematics. Cambridge University Press, 2009, vol. 3.
- [43] T. F. Gonzalez, "Clustering to minimize the maximum intercluster distance," *Theoretical Computer Science*, vol. 38, 1985.
- [44] P. Indyk, S. Mahabadi, M. Mahdian, and V. S. Mirrokni, "Composable core-sets for diversity and coverage maximization," in *Proceedings of the Symposium on Principles of Database Systems*, ser. PODS '14. New York, NY, USA: ACM Press, 2014.
- [45] J. P. S. Parkkinen, J. Hallikainen, and T. Jaaskelainen, "Characteristic spectra of munsell colors," *Journal of the Optical Society of America A*, vol. 6, no. 2, Feb. 1989.
- [46] J. Jiang, D. Liu, J. Gu, and S. Süsstrunk, "What is the space of spectral sensitivity functions for digital color cameras?" in *2013 IEEE Workshop on Applications of Computer Vision (WACV)*. New York, NY, USA: IEEE, 2013.
- [47] A. Beléndez, E. Fernández, J. Francés, and C. Neipp, "Birefringence of cellotape: Jones representation and experimental analysis," *European Journal of Physics*, vol. 31, no. 3, mar 2010.
- [48] N. Ohta, "The basis of color reproduction engineering," *Japan: Corona-sha Co*, 1997.
- [49] M. P. Johnson, "Photosynthesis," *Essays in Biochemistry*, vol. 60, no. 3, pp. 255–273, Oct. 2016.



Katherine Salesin (she/her) is a PhD student at Dartmouth College and NASA Fellow doing research in physically based rendering, advised by Wojciech Jarosz. Her research is focused on applying physically based rendering with polarization to radiative transfer problems in scientific fields such as oceanography, atmospheric science, and astronomy.



Dario Seyb (he/him) is a PhD student at Dartmouth College, focused on computer graphics, physically based rendering and real time graphics. He is particularly interested in alternative geometry and scene representations.



Sarah Friday (they/them) is a PhD student at Dartmouth College interested in research at the intersection of math, computer science, physics, and fine art. In their spare time, they practice art and serve on the Dartmouth Graduate Student Council Executive Board as the Communications Chair.



Wojciech Jarosz (he/him) is an Associate Professor of Computer Science at Dartmouth College and co-founder of the Visual Computing Lab. His group primarily focuses on physically based light transport simulation for realistic image synthesis.