

# Programmable Spectral Filter Arrays using Phase Spatial Light Modulators

Vishwanath Saragadam, Vijay Rengarajan, Ryuichi Tadano, Tuo Zhuang, Hideki Oyaizu, Jun Murayama, and Aswin C. Sankaranarayanan

**Abstract**—Computational imaging has always benefited from tools that modulate light along the many dimensions of its plenoptic function. This paper provides a practical architecture for achieving spatially varying spectral modulation using a liquid crystal phase spatial light modulator (SLM). The use of a phase SLM, however, results in strong optical aberrations due to the unintended phase modulation, thereby precluding spectral modulation at high spatial resolutions. To mitigate this, we provide a careful and systematic analysis of the aberrations arising out of phase SLMs for the purpose of spatially varying spectral modulation; this analysis results in a dual strategy of “good patterns” that minimize the optical aberrations and a deep restoration network that overcomes any residual aberrations. We show a number of unique operating points with our prototype including single- and multi-image hyperspectral imaging, material classification (fewer than two images), and dynamic spectral filtering at video rates.

**Index Terms**—Computational Imaging, Hyperspectral Imaging, Liquid Crystal, Spectral Modulation, Phase Spatial Light Modulator



## 1 INTRODUCTION

Modulating light across its many dimensions is an important enabling capability for many computational imaging and illumination systems. For example, spatial modulation of light is a critical enabler of projectors. Similarly, spatio-angular modulation of light, using a micro-lenslet array or a coded aperture, finds use in light field imaging. Hence, it is understandable that many of the advances in computational imaging have come about due to novel mechanisms for manipulating light.

We are particularly interested in techniques that can modulate the spectral content of a scene. When we look at the space of such techniques, they are largely clustered around two key properties: *programmability* that is enabled by color filter wheels, liquid crystal (LC) tunable filters [1], and dispersion-based systems [2], and *spatial selectivity*, that is enabled by assorted pixels [3, 4] where a spectral filter array is tiled on top of a sensor. There is, however, a paucity of approaches that are simultaneously *programmable* and *spatially selective*.

One approach to enable both programmability as well as spatial selectivity in spectral modulation is to use of an LC-based spatial light modulator (SLM). An LC SLM is simply a pixellated array of LC cells, where we have near-independent control over the phase retardance induced by each pixel; this retardance is converted to a spectral modulation via the use of cross polarizers. By optically aligning an image sensor to the SLM, we can leverage the degrees of freedom provided by the millions of independently-addressable pixels of the SLM to implement

a programmable spatially-varying spectral filter. Liquid crystal SLMs have been used in prior work for enhancing color gamut [5], color and polarization imaging [6] as well as hyperspectral imaging [7, 8]. However, these approaches do not address a number of practical challenges that stem from undesired properties of these SLMs. Spectral modulation with LC SLMs is invariably accompanied with an undesired phase modulation of the incident light. Further, the physical properties of LCs only allow for spatially-smooth phase retardances; hence, it is challenging to implement a diverse set of filters over any small region on the SLM. All of these artifacts are further exacerbated when we image over a large spectral range, in part due to the poor focusing performance of relay lenses.

**Contributions.** This paper proposes a system for implementing a programmable and spatially-varying spectral filter array using LC SLMs, and correcting the aberrations encountered due to non-idealities of the SLM and the associated optics. We study the factors that control undesirable phase modulation and minimize its effects by a careful design of the SLM’s retardance curve (typically, controlled as a “gamma” curve linking input 8-bit value to the path length delay induced at each pixel) as well as the patterns displayed on the SLM. To further reduce the amount of aberrations, we train a deep network to restore the spectrally-filtered measurements. We introduce a dataset, acquired using our lab prototype, consisting of images of scenes captured under a variety of spatially-varying filters. This dataset and the associated codebase can be found at [9].

We investigate three applications to showcase the effectiveness and potential of the proposed approach: hyperspectral imaging enabled by dynamically changing the bank of filters displayed on the SLM (see Fig. 1), material classification with task-specific spatially-tiled spectral filters, and the design of arbitrary spectral filters by placing the SLM in the pupil plane of the imaging system.

- Saragadam is with the Rice University, Houston, TX, USA. Corresponding author e-mail: vs44@rice.edu
- Rengarajan is with the Meta Reality Labs, Sunnyvale, CA, USA. This work was performed when he was at the Carnegie Mellon University.
- Tadano, Zhuang, Oyaizu, and Murayama are with the Sony Semiconductor Solutions Corporation, Kanagawa, Japan.
- Sankaranarayanan is with the Carnegie Mellon University, Pittsburgh, PA, USA

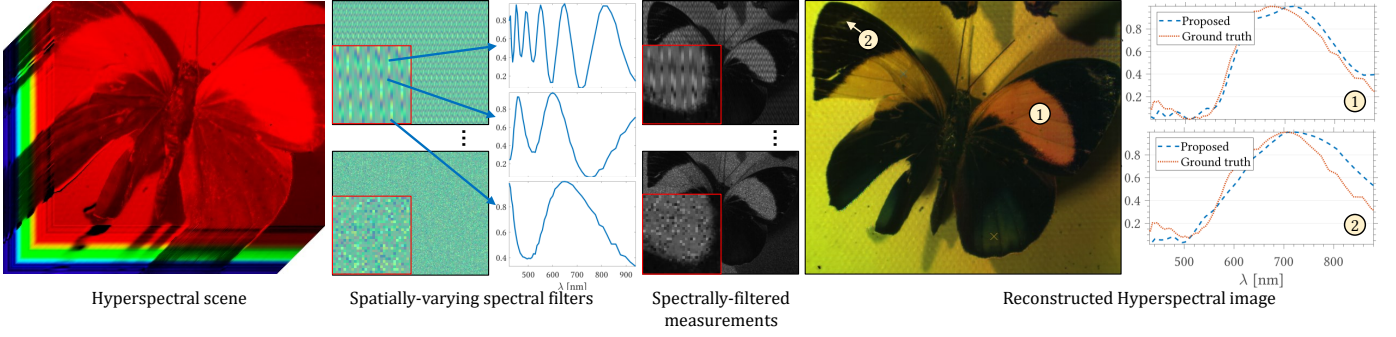


Fig. 1: **Programmable spectral filters.** We propose a practical implementation of spatially varying programmable spectral filter array. Our camera can capture intensity images under a wide configuration of spatially varying spectral filters, that enables applications that benefit from spectral measurements, namely, hyperspectral imaging, and material classification. Above, we show an example of hyperspectral imaging of a butterfly in visible to near-infrared wavelengths, reconstructed from eight images, each with a different spatial tiling of spectral filters.

**Limitations.** We inherit several limitations due to our use of phase SLMs and LC cells for spectral modulation. Spectral filters implemented using LC cells are typically restricted to sinusoidal profiles, with a spectral resolution that is inversely proportional to wavelength. Further, phase SLMs can only implement spatial patterns that are smooth, which limits the range of filters we can implement.

## 2 PRIOR WORK

We briefly discuss key results in spectral modulation, and their use in hyperspectral imaging and classification.

**Compressive hyperspectral imaging.** There have been many imaging architectures proposed that use the theory of compressive sensing to speed up HSI sensing; this includes the coded aperture snapshot spectral imager (CASSI) [10, 11], spatio-spectral coding [12], imager with DMD and a single spectrometer [13], and a conventional camera equipped with a prism in front of the lens [14]. This paper falls under the broad category of such designs.

**Spectral modulation with LC cells.** Liquid crystal cells and SLMs rely on birefringence where the two orthogonal states of polarization are delayed by varying amount applications in various fields of optics [15–17]. The spectral modulation property of LC cells have been used for compressive sensing of HSIs [18, 19] but often require several images to recover the HSI. Spectral modulation with LC cells have also been used for material classification [20]; here, a few carefully-chosen spectral filters are obtained to detect various powders in a scene. The spatially-varying modulation of SLMs has been leveraged in the past for designing color displays [5], and for spectral imaging [6–8, 8]. *However, these have not modeled or considered any of the aberrations induced in using phase SLMs for implementing spatially-varying filters; we show that the aberrations caused by the phase SLM can severely degrade reconstruction quality. To the best of our knowledge, we are the first to perform a careful analysis of this effect.*

**Assorted pixels.** Assorted pixels [3] refers to a technique where a grayscale image sensor is augmented by placing an array of filters on top to provide enhanced perception of spectrum, polarization and/or dynamic range. Assorted pixels have been extended to hyperspectral imaging as well,

where a narrowband spectral filter array is tiled on top of the sensor [4]. The main contribution of our work is in the spirit of assorted pixels; but unlike existing work, where the spectral arrays are fixed at fabrication, the proposed optical design permits a *programmable* array of spectral filters which enhances the scope of the technique in many interesting ways. For the same reasons, we refer to our approach as *programmable assorted pixels (ProAsPix)*.

## 3 PROGRAMMABLE SPECTRAL FILTER ARRAYS

We now discuss the core ideas underlying this paper, namely, spectral filtering with LC cells and the implementation of spatially-varying spectral filters using a phase SLM.

### 3.1 Basics of Spectral Filtering with LC Cells

We briefly go over the principle of operation of an LC cell when used to implement a spectral filter. The reader is referred to [1] as well as the supplemental material for a detailed treatment.

Consider an imaging setup consisting of an LC cell of thickness  $d_{LC}$  that is sandwiched between two linear cross polarizers, with their polarization axes oriented at  $\pm 45^\circ$  to the LC cell’s fast axis. Suppose that we apply a (RMS) voltage  $v$  across the LC cell which produces a birefringence  $\Delta n(v)$ . Now, *unpolarized* light incident on this setup experiences a spectral filter of the form:

$$\frac{1}{2} \left( 1 - \cos \left( 2\pi \frac{\Delta n(v) d_{LC}}{\lambda} \right) \right), \quad (1)$$

where  $\lambda$  is the wavelength of light. This filter is sinusoidal in the wavenumber, or the reciprocal of the wavelength. The frequency of this sinusoid is determined by the term  $\Delta n(v) d_{LC}$ , which quantizes the path difference introduced by the LC cell. As is to be expected, thicker LC cells introduce a large path difference which creates spectral filters with more cycles over  $\lambda$ . Similarly, higher birefringence  $\Delta n$ , that happens for low values of  $v$ , also creates a larger number of cycles over the waveband of interest.




Fig. 2: **Spectral filtering with SLMs.** (Top row) We optically align an LC-based SLM to an image sensor, using optical relays and a polarizing beamsplitter that acts as a pair of cross-polarizers before and after the SLM. This allows us to build a programmable spectral filter array, where a spatially-varying and programmable spectral filter can be realized by displaying different patterns on the SLM. (Bottom) Lab prototype.

### 3.2 Programmable Spectral Arrays using SLMs

We now describe the basic ideas underlying the proposed programmable spectral filter array. Our key insight is that an LC SLM is in essence an array of LC cells, each of which acts as a programmable spectral filter. Collocating the SLM with an image sensor, hence, allows us to obtain a device whose spectral response can be changed spatially, upto the limits imposed by the device construction. Figure 2 shows the optical schematic such a device, where the SLM and an image sensor are optically collocated to the image plane of a main lens. The polarizing beamsplitter between the SLM and the camera acts as a pair of cross-polarizers.

For sake of simplicity, let us assume that the SLM pixel size is identical to that of the sensor and that the image relay provides a one-to-one mapping between them. If we display a spatially-varying voltage pattern  $v(x, y)$  on the SLM, then the image  $i(x, y)$  observed at the sensor is given as

$$i(x, y) = \frac{1}{2} \int_{\lambda} h(x, y, \lambda) \left( 1 - \cos \left( 2\pi \frac{\Delta n(v(x, y)) d_{LC}}{\lambda} \right) \right) s(\lambda) d\lambda, \quad (2)$$

where  $s(\lambda)$  is the spectral sensitivity of the image sensor and  $h(x, y, \lambda)$  is the unmodulated hyperspectral image formed on the sensor. Hence, by appropriate choice of the pattern that we display on the SLM, we can implement spatially-varying spectral filters. The set of filters we can obtain depends on the birefringence of the SLM, the range of input voltage we can provide, and the thickness of the SLM, all of which are device specific.

### 3.3 Hardware Prototype

Figure 2 shows the lab prototype, on an optical benchtop, implementing the programmable spectral filter array. There

are some important differences between the prototype and the schematic. First, the fast and slow axes of the SLM are typically aligned to its edges; since we need the incident polarization to be at  $45^\circ$  to the fast axis, we would need to mount the SLM with an in-plane rotation. This, while feasible, is cumbersome. We instead use an achromatic quarterwave plate, immediately after the polarizing beamsplitter, with its fast axis aligned at  $45^\circ$  to the transmitted polarization state of the beamsplitter. The resulting light is circularly polarized, which has equal energy along both the fast and slow axes of the SLM, and hence, identical to rotating the SLM. Second, we use a camera lens, focused at infinity, in the relay between the SLM and the main camera; the superior multi-lens design of the camera lens has the effect of providing a dramatic improvement in the quality of captured images. Finally, we have a second image sensor, an RGB unit, placed in the previously unused arm of the beamsplitter, and collocated with the image plane of the system and, hence, the SLM and the other image sensor. We refer to this as the *guide camera*, and use it for guided filter-based reconstruction techniques that resolve the loss of spatial resolution due to tiling of the spectral filters.

**Specifications of the system.** The system is configured for a waveband spanning visible (VIS) and near-infrared (NIR) wavelengths (400 - 1000nm). We used a Holoeye Pluto2 NIR-015 SLM optimized for NIR wavebands, and endowed with a higher phase retardation which in turn provided a set of spectral filters with larger number of oscillations. The SLM resolution was  $1920 \times 1080$ , with a pixel pitch of  $8\mu m$ . The pixel pitch of the camera observing the SLM is  $6.5\mu m$ , at a resolution of  $2048 \times 2048$  pixels. The lab prototype above had a spatial field of size approximately  $10.4 \times 8$  sq.mm. We set the Fourier plane aperture to be 10mm wide; with 100mm relay lens, this resulted in an  $f/10$  aperture.

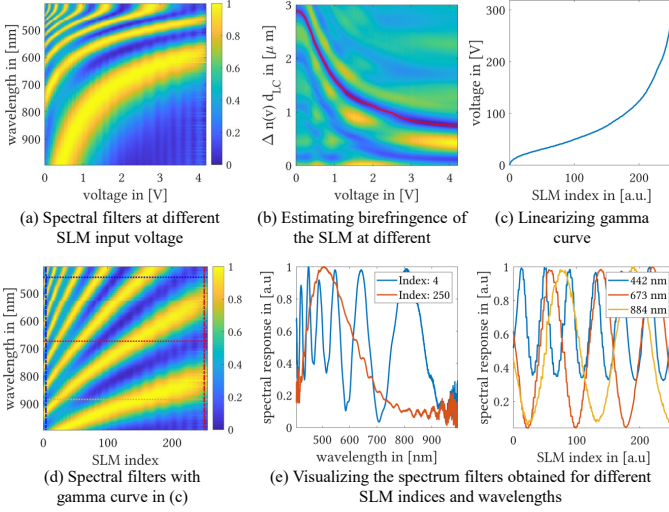
**Calibration.** We calibrated the spectral response of SLM by sweeping through 256 flat grayscaled patterns and measuring output with a spectrometer. The SLM and the spectral camera were calibrated in a pixel-to-pixel manner by sequentially scanning each row and column on the SLM. Finally, the guide camera and the spectral camera were calibrated by imaging a texture-rich scene and fitting a homography between the two images. Further details available in the supplementary.

## 4 DESIGNING SPATIALLY-VARYING FILTERS

We now discuss the choice of spatially-varying pattern that we display on the SLM. Intuitively, the pattern that we display needs to be rich over local regions, so as to provide a diverse set of spectral filters. However, the complexity of the displayed pattern also needs to be balanced against an unintended consequence of our optical system, namely, distortions to the incident wavefront due to phase modulation induced by the SLM.

**Undesired effects of phase modulation.** The proposed optical system for implementing spatially varying filters results in three key undesirable effects. First, due to non-ideal optics over broadband operation (400 - 900nm), the point spread function is spectrally varying at each spatial pixel. This problem can be addressed with complex optics (such as





**Fig. 3: Designing the SLM gamma curve.** Our goal is to estimate a gamma function that enables a linear relationship between SLM birefringence and input voltage. To do this, we use a spectrometer to measure the spectral filter produced over a range of input voltages. The measured data is shown in (a). For each voltage  $v$ , we brute-force search for the value of  $\Delta n(v)d_{LC}$  that fits the spectral filter obtained at that voltage by measuring accuracy to (1). The resulting loss function is visualized as an image in (b). Identifying the minimum for each voltage provides us with value of  $\Delta n(v)d_{LC}$ , overlaid as a red curve in (b). We use this estimated value to design a gamma curve  $\gamma(\cdot)$  for the SLM, shown in (c), such that the overall function  $\Delta n(\gamma(\cdot))d_{LC}$  is linear in its input. (d, e) The resulting set of filters that we obtain, now indexed as a function of the 8-bit index used in controlling the SLM.

microscopic tube lenses) but invariable results in extremely bulky systems. Second, the spatially varying phase modulation by the SLM distorts the wavefront of light, in the form of local tilts. This distortion is a function of the spatial gradient of the SLM pattern, and cannot be overcome even with ideal optics. Third, LC-based SLMs operate slowly. this causes the spatial patterns to be smoothed, resulting in artifacts at places of edge discontinuities. Next, we propose strategies to overcome these limitations of using an SLM for spatially varying spectral filtering.

#### 4.1 Mitigating the Effect of Phase Modulation

We first derive the expression for phase gradient and its effect on the wavefront distortion. From (1), the phase delays  $\phi(\mathbf{x}; \lambda)$ , where  $\mathbf{x} = (x, y)$ , induced by the SLM is given as

$$\phi(\mathbf{x}; \lambda) = 2\pi \frac{\Delta n(v(\mathbf{x}))d_{LC}}{\lambda}. \quad (3)$$

The local distortions to the wavefront can be characterized by the spatial gradients of  $\phi$ .

In its normal mode of operation, we do not have arbitrary control over the SLM voltage  $v(x, y)$  due to constraints on bandwidth of data on the video port used to control the SLM patterns. Instead, the voltage  $v(x, y)$  is given as

$$v(x, y) = \gamma(p(x, y)), \quad (4)$$

where  $p(x, y)$  is the 8-bit image displayed on the video port and  $\gamma(\cdot)$  is the “gamma curve” of the SLM, or the mapping of this 8-bit number to voltage applied at the SLM.

With this, the expression for the spatial gradients of  $\phi(\mathbf{x}; \lambda)$  can be written as

$$\frac{\nabla \phi}{d\mathbf{x}} = \frac{2\pi}{\lambda} d_{LC} \frac{\partial \Delta n}{d\gamma} \frac{\partial \gamma}{d\mathbf{x}} \frac{\partial p}{d\mathbf{x}} \quad (5)$$

The expression of the gradient above can be broken into three terms:

- $\frac{\partial \Delta n}{d\gamma}$ : the change in birefringence as a function of voltage across the LC cell — this is a device-specific property,
- $\frac{\partial \gamma}{d\mathbf{x}}$ : gradient of the voltage applied at an SLM pixel as a function of the SLM index value used to control the SLM; this is a term that is determined by the *gamma* curve that we select, and
- $\frac{\partial p}{d\mathbf{x}}$ : the spatial gradients of the display pattern on SLM.

Our goal is to keep the gradient as small as possible so as to minimize the undesirable effect of light tilts. While it is possible to get a zero-valued gradient simply by displaying a constant pattern  $p(x, y) = p$ , such a choice is inconsistent with a spectral filter array.

**Designing the gamma curve.** To simplify the overall design problem, while minimizing the effect of the phase gradient, we choose a gamma curve  $\gamma$  for the SLM such that

$$\frac{\partial \gamma}{\partial p} = c_0 \left[ d_{LC} \frac{\partial \Delta n}{d\gamma} \right]^{-1} \quad (6)$$

where  $c_0$  is a constant. This is equivalent to the selecting a  $\gamma$  such that  $\Delta n(\gamma(p))$  is affine in  $p$ . The procedure for designing this gamma curve requires knowledge of the phase retardance induced by the SLM at different input voltages; we illustrate the procedure for this in Figure 3 and provide a detailed description in the supplemental material. With this choice, the phase gradient in (5) reduces to

$$\frac{\nabla \phi}{d\mathbf{x}} = \frac{2\pi c_0}{\lambda} \frac{\nabla p}{d\mathbf{x}}, \quad (7)$$

and hence, the phase gradient is directly controlled by the smoothness of the pattern we display.

#### 4.2 Designing SLM Patterns

The design of the patterns displayed on the SLM needs to balance two key criteria. At one end, we need a diverse set of spectral filters in any local patch, so that we have a rich measurement operator. This richness criteria is best satisfied with patterns that are endowed with very different filters in immediate proximity. At the other end, we need to ensure that the spatial gradient of the displayed pattern is small; this is necessary so as to avoid the adverse effects of the distortions introduced by the spatially-varying phase modulation. This smoothness criteria is best satisfied by a constant pattern. Clearly, these two criteria are in direct opposition with one another. Armed with these two criteria, we consider a number of patterns, discuss the rationale behind each, and detail their relative merits. This discussion is summarized in Figure 4.

**Linear patterns and staggered variants.** We can introduce some diversity into the SLM pattern by varying it linearly in a single direction. An example of this pattern is

$$p(x, y) = x \bmod 255,$$



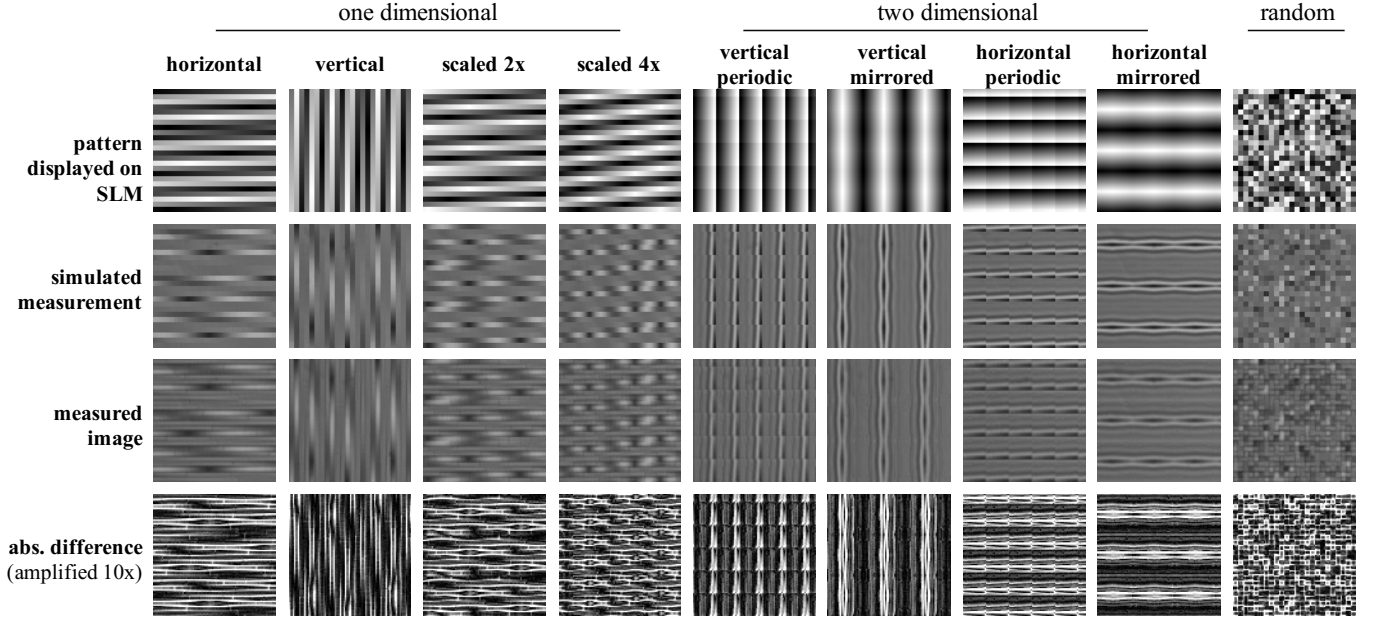


Fig. 4: **Pattern design.** We explore SLM patterns that offer varying levels of smoothness and local diversity. Each column corresponds to a different SLM pattern, spanning (columns 1-4) smooth 1D patterns, (columns 5-8) smooth 2D patterns under periodic and mirror tiling, and (column 9) random patterns. For each pattern, we show (row-1) the pattern displayed on the SLM, (row-2) a simulated image obtained from a full scan, (row-3) the measured image, and (row-4) the absolute difference between the simulated and measured image; these corresponding to imaging a (roughly) constant intensity target. Note the inaccuracies near discontinuities in the SLM image.

which has a spatial gradient of 1 intensity per SLM pixel (we will ignore these units in the sequel). This pattern is a slight improvement over the constant pattern, but has tiles that are stretched along one of the axis, which results in severe loss in spatial resolution. One approach to decrease this loss in resolution is to scale the pattern by displaying

$$p(x, y) = (2x) \bmod 255,$$

which increases the gradient by a factor of two, but decreases the tiling size by a commensurate amount. A richer way to balance out this loss of resolution is to stagger the pattern across different rows, as shown in Figure 4. This staggering intentionally introduces vertical discontinuities, whose effects can be minimized simply by repeating the pattern for a few rows and rejecting measurements at the discontinuity. Staggering allows us to avoid suffering a severe loss in resolution along one axis.

**Two-dimensional patterns.** Smooth two-dimensional patterns can be designed by tiling the 256 spectral filters into a  $16 \times 16$  tile. Here, we have multiple choices in the form of the direction of tiling, which can be horizontal, or vertical, as well as in the nature of the tiling, which can be periodic or mirror symmetric. Among these, the spatial gradients for the horizontal and vertical oriented patterns are  $240/16 \approx 15$ . Periodic tiling have strong discontinuities at the edge of each tile, but any  $16 \times 16$  tile has all the choices of filters that the SLM can offer. In contrast, symmetric tilings have no discontinuities, but do not guarantee that any  $16 \times 16$  patch covers all possible filters.

**Random patterns.** Finally, we consider random tiling of filters, which greatly increases the diversity of filters avail-

able in any local patch. To improve the smoothness of such patterns, we can repeat each random pattern in a small local window of  $3 \times 3$  or  $5 \times 5$  pixels so that the measurements at the center of each window have little aberrations.

### 4.3 Dataset

We collected a dataset of scans with our lab prototype. The dataset comprised of indoor scenes, comprising mainly of single or multiple objects (details in supplementary). We illuminated the scenes with a number of different sources including an NIR-enhanced incandescent light, a cool white LED, and CFL lamps. The acquired images from the grayscale camera and the RGB camera, are registered to the SLM using the calibration. After mapping to the SLM, the images are cropped to the central  $1024 \times 1024$  pixels. For each scene, we captured two sets of patterns that we describe next.

**Full Scan.** For each scene, we acquire a set of 256 images, corresponding to the SLM displaying a constant intensity pattern — one for each of the 8-bit intensity control that we have. Note that this results in a constant voltage for all SLM pixels, via the “gamma” function that we derive in Section 4.1. We refer to this set as the *full scan* measurements. We use it as the baseline and reconstruct a nominal ground truth hyperspectral image. This set also has none of the aberrations introduced due to phase modulation, since we are displaying a constant pattern on the SLM.

**Spatially-varying filtered images.** For each scene, we acquire a set of 92 images corresponding to the SLM displaying the patterns types described in Figure 4. For each type, we display multiple patterns by performing circular

TABLE 1: Each of the patterns in Fig. 4 under different perturbations. For the non-random patterns, these are circular shifts in one or both directions, as appropriate for the pattern. For the random pattern, we simply regenerate the random patterns. In total, we capture each scene with 92 patterns.

Name	#patterns	perturbation
1D horizontal	16	horizontal shifts
1D vertical	16	vertical shifts
1D horizontal (scale 2x)	8	horizontal shifts
1D horizontal (scale 4x)	4	horizontal shifts
2D vertical (periodic)	8	2D shifts
2D vertical (mirror)	8	2D shifts
2D horizontal (periodic)	8	2D shifts
2D horizontal (mirror)	8	2D shifts
random	16	random tiles
<b>Total</b>	<b>92</b>	

shifts of the patterns as appropriate to the type. The list of patterns with an enumeration of how we obtain a set of 92 patterns is provided in Table 1. For example, 1D patterns are shifted in just one dimension, while 2D patterns are perturbed along both dimensions. Random patterns are simply regenerated to get entirely independent images. For each of the 92 spatially-varying patterns shown on the SLM, we capture the image on both the main camera and the RGB guide camera.

#### 4.4 Artifacts in Spatially-Filtered Images

The captured spatially-varying spectral images can be compared to their ideal versions that can be assembled from the full scan measurements. Specifically, given the SLM pattern, we can create a simulated measurement where the intensity at each pixel location is the measurement in the full scan image corresponding to the SLM value at that pixel. We term these as “simulated measurements”; example of such measurements are shown in the second row of Figure 4. We immediately observe that the real measurements have a significant mismatch against the measurements simulated with the full scan data. We visualize these differences in the last row of the figure. The differences are especially significant near discontinuities of the pattern as is clearly seen in the 1D measurements, and extend beyond the discontinuities for the 2D patterns.

The mismatch between the simulated and measured images can be attributed to the following reasons. First, the SLM displays a smoothed version of the pattern we display; hence, we can expect significant mismatch between what we want to display and what the SLM implements at discontinuities. Second, there are aberrations induced due to phase distortions introduced by the SLM when we display a spatially-varying pattern; as described earlier, this leads to pattern specific errors that is prominent at places with large phase gradient. Third, the doublet lenses used in the relays introduced chromatic blur, which can be quite large given our operating range of 450-950nm, a span of wavelengths that covers visible and near-infrared bands. Due to this, a single sensor pixel measures light from multiple SLM pixels, which leads to a corrupted measurement that corresponds to a blurred measurement in both space as




Fig. 5: **Restoration network architecture.** Our pattern-independent decoder architecture rectifies the distortions in the measurement image using positional codes based on coordinates and pattern index.

well as the spectrum. Note that this chromatic blur is also present in the full scan data; however, since the full scan data is measured with a constant spectral filter, the resulting blur is spatially *invariant*. To reduce these non-idealities in measured data, we rely on a learning-based approach using neural networks.

#### 4.5 Deep Restoration of Spatially-Filtered Images

Our goal is to learn a mapping that takes the distorted measurements and produce clean “simulated” measurements. Since interactions between neighboring spatial filters are non-linear, and often spatially varying, simple models may not capture the non-idealities. We instead rely on a learning-based approach to produce accurate measurements.

We design a single neural network that takes the measured image as input, for any spatially-varying pattern on the SLM that is also provided as input, and outputs the “simulated” measurement for that pattern produced from the full scan data. To account for the spatially varying artifacts, we include the coordinates  $(x, y)$ , the pattern index  $p$  to the inputs, in addition to the measured intensity at each pixel. We leverage recent advances in positional encoding [21, 22] to build a 64-length input vector at each pixel. Thus, the input to the network consists of a measurement intensity channel, and 64 channels each for  $x, y$  and  $p$ , resulting in a total of 193 channels, as shown in Figure 5. The output of the network is a single restored image channel corresponding to the simulated measurement.

**Network Architecture.** We use an encoder-decoder architecture similar to that of U-net proposed by [23] with four downsampling and four upsampling blocks as shown in Figure 5. The specifics of this network and how it was trained is provided in the supplement.

**Restoration.** During inference, we input the measured image along with the positional encoding channels (all as  $1024 \times 1024$  images) to the trained network to get the restored image as output which will be used for hyperspectral image reconstruction. We show the results of our restoration network for a test scene which is not used in training and validation in Figure 6.

**Restoration of unseen patterns.** Since our restoration network is trained with the pattern code as well as an input, it is capable of restoring unseen patterns as well. We show the restoration outputs of patterns, previously unseen during training, in Figure 7. Our network successfully restores these patterns, albeit with a slight loss in performance as compared to the patterns used at training. This indicates that




Fig. 6: **Restoration of measured image.** We show  $128 \times 128$  patches of a scene from our test data. Each row in the left and right page columns corresponds to some of the different patterns that we employed during capture. The simulated measurement in (a) is the “ground-truth” image for each pattern. The measured image in (b) is very much distorted as can be seen by its absolute difference with the simulated measurement in (c). Our restoration network cleans up the distortions and produces the restored measurement image shown in (d) which has much lower distortions as can be seen in (e). The efficacy of our restoration network can also be seen from much higher reconstruction PSNR values which are shown as insets.

our trained network does not merely memorize a pseudo-random distribution of the patterns present in the training set, but can generalize for any pattern.

## 5 APPLICATION: HYPERSPECTRAL IMAGING

We now consider the application of hyperspectral imaging (HSI) with phase SLMs; our eventual goal is not just to provide a viable framework for HSI but also to quantify the improvement provided by the deep restoration framework.

The image measurements made in (2) can be written as,

$$i_k(x, y) = \int_{\lambda} h(x, y, \lambda) f_{\text{SLM}}^k(x, y, \lambda) d\lambda, \quad (8)$$

where  $f_{\text{SLM}}^k$  is an instance of the spatially-varying spectral response of the SLM. Discretizing and converting images to vectors we get  $\mathbf{i}_k = X\Phi_k$ , where  $X \in \mathbb{R}^{N_x N_y \times N_{\lambda}}$  is the matrix representation of the HSI of the scene. We repurpose the body of work devoted to HSI reconstruction to solve (8) to obtain  $h(x, y, \lambda)$  from  $\{\mathbf{i}_k(x, y)\}$ .

**Guide-free reconstruction.** The traditional approach is to formulate the recovery as a linear inverse problem,

$$\min_X \sum_{k=1}^N \|\mathbf{i}_k - X\Phi_k\|^2 + \mathcal{R}(X), \quad (9)$$

where  $\mathcal{R}(\cdot)$  is a spatial/spectral regularizer. For snapshot approaches ( $N = 1$ ), we used an untrained deep network as a regularizer, similar to the deep image prior (DIP) [24]. We expressed the HSI as the output of a convolutional neural network (CNN) equipped with 2D spatial convolutions and whose input was a fixed noise pattern. We then solve for




Fig. 7: **Restoration of measured image for unseen patterns.** Our trained network successfully cleans up the distortions in the measured image even for completely unseen patterns such as for the checkerboard-like patterns shown here.

the HSI by optimizing for the CNN’s parameters. For multi-pattern reconstruction ( $N \geq 2$ ), we found that a combination of 2D total variation (TV) prior and a 1D spectral smoothness prior enabled high quality reconstruction.

**Guided reconstruction.** In the presence of an RGB guide image, we leverage a super pixelation-based reconstruction technique inspired by recent works in hyperspectral imaging [25]. Given an RGB image  $I_{\text{RGB}}[x, y]$  of the scene, we first partition the image into  $Q$  superpixels [29]. We make the assumption that the spectral profiles associated with pixels within any superpixel are scaled multiples of each other. Then we model the HSI within each super pixel as a rank-1 matrix, where the spatial component is a scaled version of the grayscale image. Specifically,  $X_q = \mathbf{g}_q \mathbf{s}_q^T$  where  $\mathbf{g}_q$  is the grayscale image intensity in the  $q^{\text{th}}$  super pixel and  $\mathbf{s}_q$  is the spectral basis in the  $q^{\text{th}}$  super pixel. We then solve for a local least squares problem to estimate the spectral component of the rank-1 matrix,

$$\min_{\mathbf{s}_q} \sum_k \|\mathbf{i}_{k,q} - \mathbf{g}_q \mathbf{s}_q^T \Phi_{k,q}\|^2 + \eta \|\mathbf{s}_q\|^2. \quad (10)$$

Details about reconstruction in supplementary material.

### 5.1 Simulation

Figures 8 and 9 provide reconstruction results on simulated data with single and multiple spectrally coded measurements. Specifically, we simulated the acquisition setup on several datasets including the ICVL [27] and KAIST [26] datasets consisting of hyperspectral images over visible wavelengths. We compare our technique against existing snapshot techniques including Choi et al. [26] which consists of a CASSI-type hardware and a deep neural network based reconstruction, and SASSI [25] which consists of a sparse spatio-spectral sampler along with RGB fusion. We also compared to two techniques that recovered HSIs from RGB images [27, 28]. We simulated photon noise by assuming a maximum light level of 1000 photoelectrons, and a readout noise of 2 electrons, resulting in a signal to noise ratio of 30dB. For ProAsPix, we used the symmetric 2D horizontal pattern with for all our simulations. We quantify performance using peak signal to noise ratio (PSNR) metric and spectral angular mapping (SAM). Further details about optimization are in the supplementary.



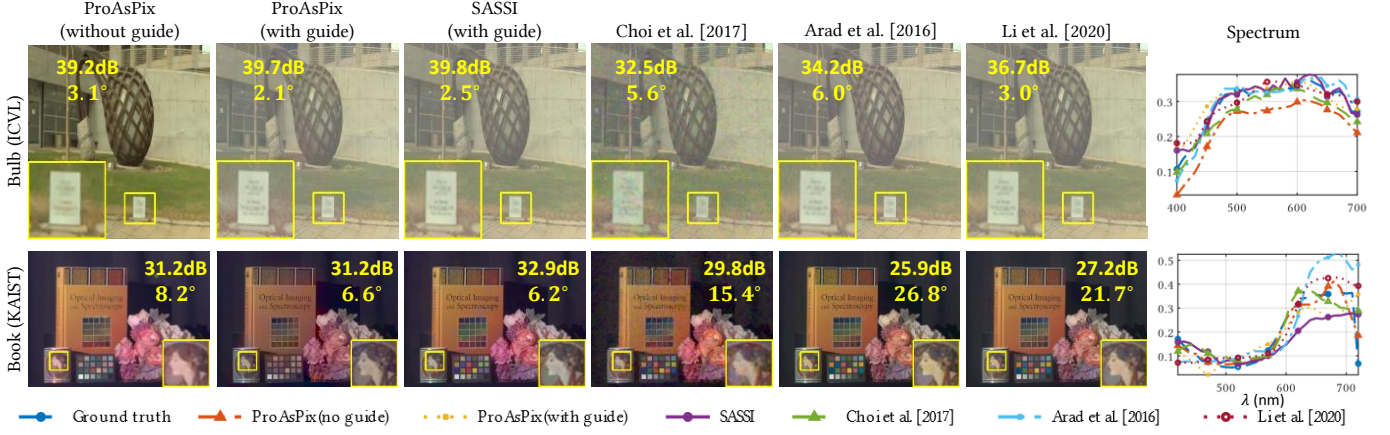


Fig. 8: **Comparisons against snapshot approaches in visible wavelengths.** We compare ProAsPix, with and without guide, to a number of single-image competitors including spatial coding techniques—SASSI [25] and Choi et al. [26]—as well as RGB to HSI techniques [27, 28]. Across the board, ProAsPix with and without guide images performs comparable to SASSI and outperforms all other approaches, particularly techniques that recover hyperspectral images from RGB images.

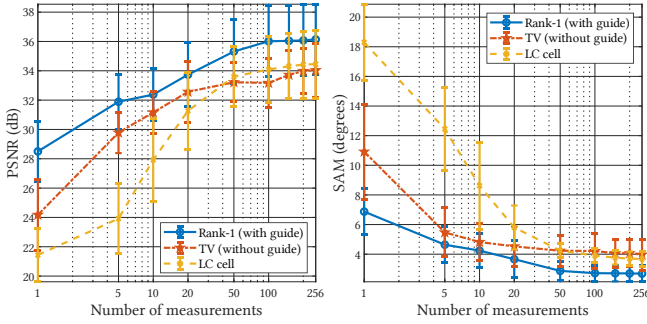


Fig. 9: **Multi-frame performance.** Our proposed approach, with and without a guide image, outperforms the spatially-constant modulation of LC cells, especially when imaging with a small number of image measurements.

**Snapshot reconstruction in the visible domain.** Most HSI reconstruction techniques based on learned models are fine-tuned for the visible domain. Figure 8 shows reconstruction for HSIs in the visible wavelengths (400 - 700nm) which included SASSI [25] (sparse spatio-spectral + guide image), [26] (SD-CASSI + learned reconstruction), [27] (dictionary-based RGB to HSI), and [28] (learned RGB to HSI). Across the board, ProAsPix results were qualitatively and quantitatively superior. In part, this can be attributed to our approach not multiplexing spatial and spectral dimensions as is commonly done in CASSI and its variants. RGB to HSI techniques [27, 28] were trained on the ICVL dataset; hence, the results for the HSI from this dataset (top row) were comparable to ProAsPix.

**Comparisons with multiple captures.** Our primary competitor for multi frame reconstruction is a single LC cell capture [18] where images are captured with a spatially *invariant* spectral modulation. Figure 9 plots reconstruction accuracy with LC cell and our technique for varying number of images. ProAsPix with guided filtering uniformly outperforms the spatially-invariant LC cell. The reconstruction in the absence of a guide image is better than LC cell at fewer measurements and similar with 50 or more images. For a small number of images, it is more advantageous to

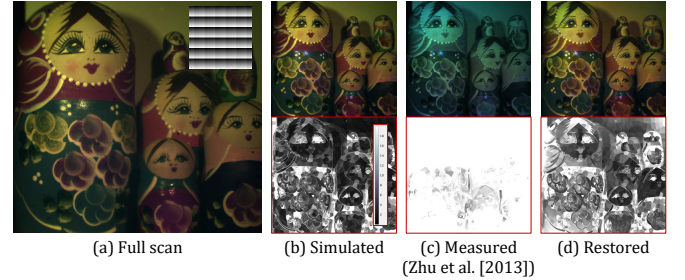


Fig. 10: **Comparison of simulated, measured, and restored measurements.** (a) Shows an RGB image of the scene with inset of the spectral filter array that we used for the results in the other columns. (b-d) visualizes reconstructed HSIs using the simulated, measured ([7]), and restored measurements as rendered RGB images. Below each reconstruction, we show the angular error against the full scan reconstructions; for these error maps, the brightest values are errors that are  $20^\circ$  or higher.

spatially multiplex the various spectral filters. In all, our approach can achieve 28dB (SAM of 10 degrees) or better with even a single image, and surpasses 34dB (SAM of 4 degrees) with as few as 20 images. We also observe that with a large number of images, spatial multiplexing has a similar effect to capturing images with spatially invariant spectral filters. Additional results in supplementary material.

## 5.2 Real Results

**Setup.** We recover HSIs at a spatial resolution of  $1024 \times 1024$  and a spectral resolution of 53 bands in the span of 420 to 940 nm. Since the spectral filters are linear in  $1/\lambda$ , we sample these 53 bands uniformly in the reciprocal of the wavelength, which provided a small performance increase over linear sampling in wavelength.

**HSI reconstruction from a single image.** Figure 10 provides an example of the reconstructions obtained with the simulated, measured, and restored measurements. Evidently, accounting for the aberrations dramatically improves the performances and produces results similar to ideal simulated conditions. Next, we observe that 2D patterns gen-




Fig. 11: **Patterns used for multi-image hyperspectral imaging.** We greedily chose up to 16 patterns that ensured diversity of spectral filters at each pixel.

erally lead to better reconstructions in real experiments too, for the simulated as well as the restored measurements. This indicates that the local diversity of patterns does play an important role in reconstruction performance.

**Multi-image reconstructions.** For multi-image reconstructions, we greedily selected a sequence of patterns that provide maximal diversification of filters. Starting with the best performing pattern shown in Figure 10, which is a 2D horizontal/periodic pattern, we sequentially add patterns that maximizes the addition of new spectral pixel over all previously selected patterns. The sixteen such patterns that we select with this scheme is visualized in Figure 11. We compare the reconstructions with rank-1 guided filter and guide-free TV prior in Figure 12. For the rank-1 approach, we linearly increase the number of superpixels used with the number of measurements. We also compare to reconstructions that would be obtained with just an LC cell, as opposed to an SLM, which would only provide global spectral modulation. Overall, ProAsPix works significantly better than what we get with an LC cell.

## 6 OTHER APPLICATIONS

### 6.1 Spatial Tiling for Material Classification

A powerful capability enabled by our system is that of disambiguating between different materials in a scene. To obtain an optimal filter array, given a selection of  $K$  materials, we first measure their response to the 256 spectral filters that can be created by the SLM. An example of these measurement traces can be seen in Figure 13(b). We can now select a few, typically a small number of filters, say  $Q$  of them, such that measurements associated each of the  $K$  materials are distinct. We perform a brute-force / greedy scan to identify a set of  $Q$  filters that produce the maximal-minimum distance of features on the simplex. Now that we have our  $Q$  filters, we tile them spatially in as compact a block as possible. For example, in Figure 13(c), we tile  $Q = 3$  filters in  $2 \times 2$  blocks by repeating one of them to differentiate between real and fake plant. Once we capture a measurement image, we first demosaick the measurement using standard linear filtering to get a  $Q$ -channel image. We project each pixel onto the simplex, and perform nearest neighbor classification on the resulting feature to obtain the material map of the scene (Fig. 13(e)). This result indicates the immense potential of our system for applying adaptive sensing techniques on top of the setup.

### 6.2 Application: Arbitrary Spectral Filters

An important capability in spectral filtering is the ability to have a programmable filter capable of displaying an arbitrary

shape. LC cells can only implement sinusoidal spectral filters; other results have shown that stacking multiple LC cells offers the ability to get narrowband filters [16, 17]. We explore implementation of spectral filters with profiles that go beyond sinusoids by placing the SLM in the pupil plane, allowing the SLM to modulate light from all scene points. However, placing the SLM in the Fourier plane results in spatially-invariant filtering in the image plane. A key advantage of using SLM is that the switching time between any two filters is same as display rate (60Hz). In contrast, LC tunable filters require hundreds of milliseconds to switch between two wavelengths, which gets larger with increasing distance between the two wavelengths.

Suppose that we seek to implement the spectral filter  $\mathbf{s} \in \mathbb{R}^{N_\lambda}$ . Let  $\Lambda \in \mathbb{R}^{N_\lambda \times 256}$  be the matrix of 256 spectral responses corresponding to different SLM voltages. Our goal is to solve for weights  $\mathbf{w} \in \mathbb{R}^{256}$  such that  $\mathbf{s} = \Lambda \mathbf{w}$ ; here, the  $i$ -th element of  $\mathbf{w}$ , which we denote of  $w_i$  provides the contribution from the corresponding filter. Once we have  $\mathbf{w}$ , we design an SLM pattern which allocates an area corresponding to  $w_i$  to the  $i$ -th index value.

**Estimating the weight vector  $\mathbf{w}$ .** It is also important the  $\mathbf{w}$  be constrained to be positive since we cannot allocate a negative area in the SLM plane to a filter. We can estimate such positive patterns by either solving a nonnegative least squares problem, or by solving an unrestricted pattern, and displaying positive and negative parts of the filter separately, capturing two images and subtracting them. The latter would require an additional image, but provides a greater space of filters that we can potentially implement.

**Implementing the weight vector on the SLM.** Given a positive weight vector  $\mathbf{w}$ , implementing it on the SLM requires us to divide the SLM area into multiple regions such that the  $i$ -th region displays the SLM index  $i$  and has an area proportional to  $w_i$ . We rely on a simple heuristic for performing this allocation. Specifically, SLM pixels in the same column are forced to have the same value, and so we allocate each of the 1024 columns at our disposal to the 256 values proportional to the values in  $\mathbf{w}$ . Figure 14 shows simulation and real results of implementing bandpass filters. We generated a set of Gaussian filters with varying center wavelength and a variance of  $30^2 \text{ (nm)}^2$ , as well as edge filters. For each spectral profile, we implemented both standard and non-negative least squares, which are both plotted on top of the ground truth profile. We can observe that non-negative least square profiles are less precise than that of least squares, for aforementioned reasons. We also implemented the non-negative filter weights in the lab prototype with a spectrometer in the imaging arm; the results from which are shown in the figure as well. We observe that non-negative least squares follows the simulations results for the most part. Least squares offers significantly better fit, both in simulation and in real capture but suffers from photon noise. The real spectral profiles oscillate more significantly at the deep blue and far red ends, likely due to poor efficiency of SLM and optics, and low intensity of the incandescent lamp at those wavelengths.

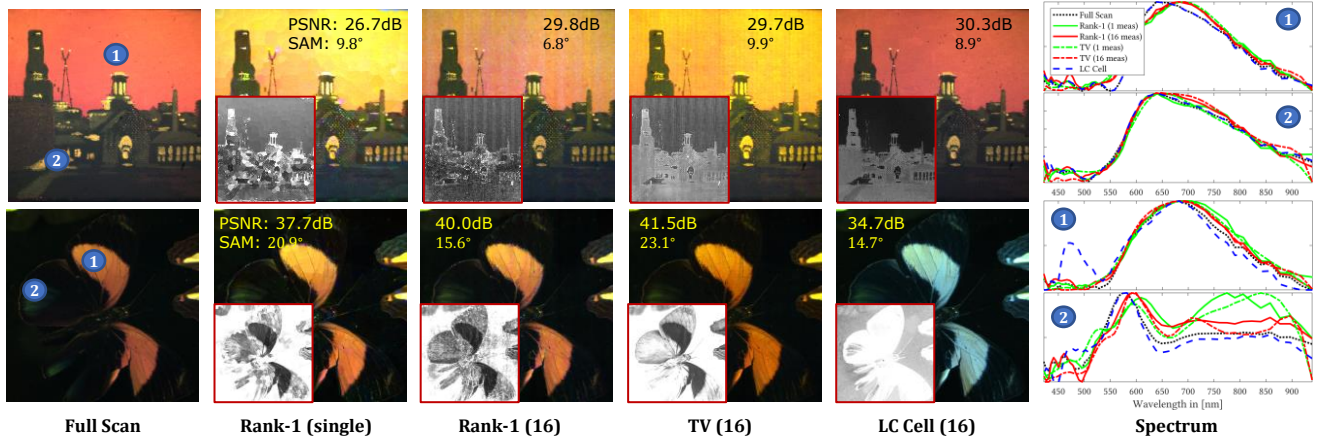


Fig. 12: **Qualitative evaluation for multiple patterns reconstructions.** We compare reconstructions from full scan of 256 measurements, a single image reconstruction using the guided Rank-1 technique, and 16 image reconstructions from both the guided rank-1 and guide-free TV techniques. We compare these to reconstructions from an LC cell, with 16 measurements as well. For each scene, we show rendered RGB images and spectrum at two points marked in the first column. Inset on the RGB images are the angular errors, visualized as in other figures with a range of 0 to 20°.

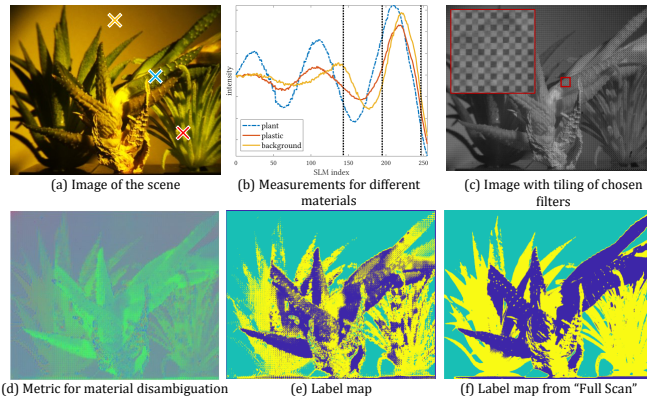


Fig. 13: **Disambiguating between materials using programmable spectral filter arrays.** (a) We image a scene with plants, real and plastic. (b) The measurement trace, as a function of SLM input index, is visualized for the two materials as well as the background. We find three index values, marked with dotted vertical lines, that lead to maximally different measurements for the three materials. (c) An image of the scene is captured with an SLM displaying a checkerboard pattern comprising of the two chosen index values. The inset is the zoomed in version of the cropped region marked in red. (d) Using this single measurement, we can now create a metric that maximally disambiguates between the two materials and threshold it in (e) to get a material map. For comparison, the label map from the full scan is shown in (f).

## 7 DISCUSSIONS

This paper introduces a novel technique for spectral modulation — namely, a programmable and spatially-varying spectral filter array — and discusses its use in single- and multi-shot hyperspectral imaging. We achieve this capability using an LC-based phase SLM, and develop an optical schematic for implementing it while computationally handling unmodeled aberrations in the setup.

**Enhancing spectral diversity of filters.** The richness of the spectral modulation produced by our system relies on the range of phase retardation that can be implemented by the




Fig. 14: **Implementing arbitrary spectral filters.** The figure shows band pass Gaussian filters (top row) with a standard deviation of 30 nm, and edge filters (bottom row) with different center frequency. For each target filter (black), we estimate SLM spatial pattern using unconstrained (green) and non-negative least squares fits (red). The plots show both simulated results that provide ideal implementation conditions for this SLM, and real captures using a lab spectrometer. Non-negative least squares provides a poor fit, both in simulation and hardware, likely due to the sinusoidal filter profiles constraining the space of implementable non-negative filters. Unconstrained least square does a significantly better job but is susceptible to noise, especially at deep blue and far red wavelengths, where our light source does not have sufficient spectral radiance.

SLM. For our system, this range spans  $3\mu\text{m}$  to  $800\text{nm}$  — increasing this range is an important direction in enhancing the utility of our design. One way of realizing this is by introducing an LC cell in front of the SLM and using the additional phase retardation provided by it. Figure 15 visualizes the range of spectral filters we can obtain once we add such an LC retarder (Thorlab LCC1115-B) immediately in front of the SLM. The resulting setup implements spectral






Fig. 15: **Increasing range of spectral modulation using an LC cell in the optical pathway.** We can use an additional LC cell to increase the overall phase retardance, and hence the diversity of spectral filters implemented by our system. Shown are the filters implemented by the SLM at two different input voltages across the LC cell.

filters that have the form

$$\frac{1}{2} - \frac{1}{2} \cos \left( 2\pi \frac{\Delta n(v_{\text{SLM}}(x, y))d_{\text{SLM}} + 2n(v_{\text{LC}})d_{\text{LC}}}{\lambda} \right), \quad (11)$$

where the terms marked with “SLM” and “LC” denote to retardances applied at the SLM and LC cell, respectively. We observe that controlling the voltage across this retarder shifts the range of phase retardance that we can apply with our setup. Since the LC cell has the largest birefringence when no voltage is applied across it, we see the spectral filters with the largest oscillations at low voltages; however, note that the diversity of filters is low in this setting since the added phase retardance overwhelms the range of the SLM. At a slightly higher voltage (1.5V in Figure 15), we observe a different set of filters with greater diversity but fewer cycles. Hence, adding an LC cell opens up a novel and richer design space and is likely a powerful addendum to our design.

**Reconstruction algorithms.** Recent work on learning-based approaches have provided significant improvements in hyperspectral image reconstruction. Example of this include Choi et al. [26] for the CASSI architecture and Gedalin et al. [30] for LC cell-based imaging. These advances are largely complimentary to the optical setup introduced in this paper, i.e., we can expect the quality of reconstructions obtained with our setup to improve with the use of such sophisticated reconstruction algorithms.

We derived a set of patterns that are well-suited for a large spectrum of applications. However, our approach will benefit from learning-based approaches for pattern selection. End-to-end techniques that simultaneously optimize the pattern along with reconstruction are expected to increase the quality of results and will be pursued as future work.

**Miniaturization.** Our current setup involves 2 optical relays which increases the bulk as well as the spatial and chromatic aberrations in the system. All of these concerns are ameliorated if we used a transmissive SLM instead of a reflective one, which would allow us to place the SLM directly on top of the sensor. This design would likely become a viable alternative to many existing hyperspectral imaging techniques, once we have such transmissive SLMs with larger phase retardance and smaller pixel pitch.

## ACKNOWLEDGEMENTS

This work was support in part by the National Geospatial-Intelligence Agency’s Academic Research Program (Grant #: HM0476-22-1-0004), a Sony Faculty Innovation Award, and the National Science Foundation (Grants #: 1652569, 1730147, 1801382).

## REFERENCES

- [1] S.-T. Wu, U. Efron, and L. D. Hess, “Birefringence measurements of liquid crystals,” *Appl. Opt.*, vol. 23, no. 21, pp. 3911–3915, Nov 1984.
- [2] A. Mohan, R. Raskar, and J. Tumblin, “Agile spectrum imaging: Programmable wavelength modulation for cameras and projectors,” in *Comp. Graphics Forum*, 2008.
- [3] S. G. Narasimhan and S. K. Nayar, “Enhancing resolution along multiple imaging dimensions using assorted pixels,” *IEEE Trans. Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 27, no. 4, pp. 518–530, 2005.
- [4] A. Lambrechts, P. Gonzalez, B. Geelen, P. Soussan, K. Tack, and M. Jayapala, “A cmos-compatible, integrated approach to hyper- and multispectral imaging,” in *IEEE International Electron Devices Meeting*, 2014.
- [5] W. Harm, A. Jesacher, G. Thalhammer, S. Bernet, and M. Ritsch-Marte, “How to use a phase-only spatial light modulator as a color display,” *Opt. letters*, vol. 40, no. 4, pp. 581–584, 2015.
- [6] T.-H. Tsai, X. Yuan, and D. J. Brady, “Spatial light modulator based color polarization imaging,” *Opt. Express*, vol. 23, no. 9, pp. 11 912–11 926, May 2015.
- [7] R. Zhu, T.-H. Tsai, and D. J. Brady, “Coded aperture snapshot spectral imager based on liquid crystal spatial light modulator,” in *Frontiers in Optics*, 2013.
- [8] J. Chen, M. Hirsch, B. Eberhardt, and H. Lensch, “A computational camera with programmable optics for snapshot high-resolution multispectral imaging,” in *Asian Conf. Computer Vision*, 2018.
- [9] “Code for programmable spectral filter arrays,” <https://github.com/Image-Science-Lab-cmu/ProAsPix>, 2023.
- [10] M. Gehm, R. John, D. Brady, R. Willett, and T. Schulz, “Single-shot compressive spectral imaging with a dual-disperser architecture,” *Opt. Express*, vol. 15, no. 21, pp. 14 013–14 027, 2007.
- [11] A. Wagadarikar, R. John, R. Willett, and D. Brady, “Single disperser design for coded aperture snapshot spectral imaging,” *Appl. Optics*, vol. 47, no. 10, pp. B44–B51, 2008.
- [12] X. Lin, G. Wetzstein, Y. Liu, and Q. Dai, “Dual-coded compressive hyperspectral imaging,” *Optics Letters*, vol. 39, no. 7, pp. 2044–2047, 2014.
- [13] T. Sun and K. Kelly, “Compressive sensing hyperspectral imager,” in *COSI*, 2009.
- [14] S.-H. Baek, I. Kim, D. Gutierrez, and M. H. Kim, “Compact single-shot hyperspectral imaging using a prism,” *ACM Transactions on Graphics (TOG)*, vol. 36, no. 6, pp. 1–12, 2017.
- [15] J. Beeckman, K. Neyts, and P. J. Vanbrabant, “Liquid-crystal photonic applications,” *Optical Engineering*, vol. 50, no. 8, p. 081202, 2011.
- [16] B. Lyot, “Optical apparatus with wide field using interference of polarized light,” *CR Acad. Sci.*, vol. 197, no. 1593, 1933.
- [17] Y. Öhman, “A new monochromator,” *Nature*, vol. 141, no. 3563, pp. 291–291, 1938.
- [18] Y. Oiknine, I. August, V. Farber, D. Gedalin, and A. Stern, “Compressive sensing hyperspectral imaging by spectral multiplexing with liquid crystal,” *Journal of Imaging*, vol. 5, no. 1, p. 3, 2019.

- [19] Y. August and A. Stern, "Compressive sensing spectrometry based on liquid crystal devices," *Optics letters*, vol. 38, no. 23, pp. 4996–4999, 2013.
- [20] T. Zhi, B. R. Pires, M. Hebert, and S. G. Narasimhan, "Multispectral imaging for fine-grained recognition of powders on complex backgrounds," in *IEEE Intl. Conf. Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [21] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," 2017.
- [22] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "NeRF: Representing scenes as neural radiance fields for view synthesis," in *European Conf. Computer Vision (ECCV)*, 2020.
- [23] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *MICCAI*, 2015, pp. 234–241.
- [24] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Deep image prior," in *IEEE Intl. Conf. Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [25] V. Saragadam, M. De Zeeuw, R. G. Baraniuk, A. Veeraraghavan, and A. C. Sankaranarayanan, "SASSI — Superpixelated adaptive spatio-spectral imaging," *IEEE Trans. Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 43, no. 07, pp. 2233–2244, July 2021.
- [26] I. Choi, D. S. Jeon, G. Nam, D. Gutierrez, and M. H. Kim, "High-quality hyperspectral reconstruction using a spectral prior," *ACM Trans. Graphics*, vol. 36, no. 6, pp. 218:1–13, 2017.
- [27] B. Arad and O. Ben-Shahar, "Sparse recovery of hyperspectral signal from natural rgb images," in *European Conf. Computer Vision (ECCV)*, 2016.
- [28] J. Li, C. Wu, R. Song, Y. Li, and F. Liu, "Adaptive weighted network with camera spectral sensitivity prior for spectral reconstruction from rgb images," in *IEEE Intl. Conf. Computer Vision and Pattern Recognition (CVPR) workshops*, June 2020.
- [29] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk, "Slic superpixels compared to state-of-the-art superpixel methods," *IEEE Trans. Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 34, no. 11, pp. 2274–2282, 2012.
- [30] D. Gedalin, Y. Oiknine, and A. Stern, "DeepCubeNet: Reconstruction of spectrally compressive sensed hyperspectral images with deep neural networks," *Opt. Express*, vol. 27, no. 24, pp. 35 811–35 822, 2019.



**Vishwanath Saragadam** is a postdoctoral researcher at Rice University with Prof. Richard G. Baraniuk and Prof. Ashok Veeraraghavan. Vishwanath's research is at the intersection of computational imaging, meta-optics, and neural representations, and focuses on co-designing optics, sensors, and algorithms for solving challenging problems in vision. He received his PhD from Carnegie Mellon University, advised by Prof. Aswin Sankaranarayanan, where his thesis won the A. G. Jordan outstanding thesis award in

2020. Vishwanath is a recipient of the outstanding postdoctoral research award at Rice in 2023, the best paper award at ICCP 2022, Prabhu and Poonam Goel graduate fellowship in 2019, and an outstanding teaching assistant award in 2018.



**Vijay Rengarajan** is a research scientist at Meta Reality Labs. His research interests are in the intersection of computational photography and image processing. He received his PhD from the Indian Institute of Technology, Madras in 2017, following which he was a postdoctoral researcher at the Carnegie Mellon University.



**Ryuichi Tadano** received the B.E. degree in information and communication engineering and the M.S. degree in frontier informatics from The University of Tokyo, Tokyo, Japan, in 2005 and 2007, respectively. He is currently a Senior Manager at Sony Semiconductor Solutions Corporation, Kanagawa, Japan. He was a Visiting Scholar with Rice University, Houston, TX, USA, from 2014 to 2015. His research interests lie in the areas of computer vision, signal processing, and computational imaging.



dustrial applications, agriculture sensing.

**Tuo Zhuang** received the B.E. degree in information technology from Nagoya Institute of Technology and M.S. degree in information science from Nara Institute of Science and Technology. He was a Research Engineer in Omron Corporation since 2006. He is currently a Research Engineer at Sony Semiconductor Solutions Corporation, Kanagawa, Japan. Since he joined Sony Corporation in 2016, he has been developing computer vision systems, optical systems, and image processing algorithms for industrial applications, agriculture sensing.



**Hideki Oyaizu** received the M.S. degree in information technology from Tohoku University. He is currently a Research Engineer at Sony Semiconductor Solutions Corporation, Kanagawa, Japan. Since he joined Sony Corporation, in 1989, he has been developing visual inspection systems, image processing algorithms, stereo camera systems, and video-processing algorithms for industrial applications, gaming, and sensing.



**Jun Murayama** received the BSc degree Electrical Engineering from Kyoto University in 1991. He is currently a Chief Image Processing Researcher at Sony Semiconductor Solutions Corporation, Kanagawa, Japan. His research interests include computational photography, image processing and novel sensing devices.



**Aswin C. Sankaranarayanan** is a professor in the ECE department at CMU, where he is the PI of the Image Science Lab. His research interests are broadly in compressive sensing, computational photography, signal processing and machine vision. His doctoral research was in the University of Maryland where his dissertation won the distinguished dissertation award from the ECE department in 2009. Aswin is the recipient of best paper awards at SIGGRAPH 2023, CVPR 2019 and ICCP 2021 & 2022, the CIT Dean's Early Career Fellowship, the Spira Teaching award, the NSF CAREER award, the Eta Kappa Nu (CMU Chapter) Excellence in Teaching award, and the Herschel Rich Invention award from Rice University.