

Single-shot compressive spectral imaging with a dual-disperser architecture

M. E. Gehm¹, R. John², D. J. Brady^{2*}, R. M. Willett², and T. J. Schulz³

¹University of Arizona Electrical and Computer Engineering Department, Box 210104,
Tucson, Arizona 85721

²Duke University Electrical and Computer Engineering Department, Box 90291, Durham,
North Carolina 27708

³Michigan Technological University Electrical Engineering Department, 1400 Townsend Dr.,
Houghton, Michigan 49931

*Corresponding author: dbrady@ece.duke.edu

Abstract: This paper describes a single-shot spectral imaging approach based on the concept of compressive sensing. The primary features of the system design are two dispersive elements, arranged in opposition and surrounding a binary-valued aperture code. In contrast to thin-film approaches to spectral filtering, this structure results in easily-controllable, spatially-varying, spectral filter functions with narrow features. Measurement of the input scene through these filters is equivalent to projective measurement in the spectral domain, and hence can be treated with the compressive sensing frameworks recently developed by a number of groups. We present a reconstruction framework and demonstrate its application to experimental data.

© 2007 Optical Society of America

OCIS codes: (110.0110) Imaging systems; (120.6200) Spectrometers and spectroscopic instrumentation

References and links

1. P. A. Townsend, J. R. Foster, J. R. A. Chastain, and W. S. Currie, "Application of imaging spectroscopy to mapping canopy nitrogen in the forests of the central Appalachian Mountains using Hyperion and AVIRIS," *IEEE Trans. Geosci. Remote Sens.* **41**, 1347–1354 (2003).
2. W. Smith, D. Zhou, F. Harrison, H. Revercomb, A. Larar, A. Huang, and B. Huang, "Hyperspectral remote sensing of atmospheric profiles from satellites and aircraft," in *Proc. SPIE*, vol. 4151, pp. 94 – 102 (2001).
3. R. P. Lin, B. R. Dennis, and A. O. B. (Eds.), *The Reuven Ramaty High-Energy Solar Spectroscopic Imager (RHESSI) - Mission Description and Early Results* (Kluwer Academic Publishers, Dordrecht, 2003).
4. W. Denk, J. Strickler, and W. Webb, "Two-photon laser scanning fluorescence microscopy," *Science* **248**(4951), 73–76 (1990).
5. R. Schultz, T. Nielsen, J. Zavaleta, R. Ruch, R. Wyatt, and H. Garner, "Hyperspectral imaging: A novel approach for microscopic analysis," *Cytometry* **43**, 239 – 247 (2001).
6. M. Hinnrichs, J. Jensen, and G. McAnally, "Handheld hyperspectral imager for standoff detection of chemical and biological aerosols," *Proc. SPIE* **5268**, 67–78 (2003).
7. J. Mooney, V. Vickers, M. An, and A. Brodzik, "High-throughput hyperspectral infrared camera," *J. Opt. Soc. Am. A* **14**, 2951–2961 (1997).
8. M. Descour, C. Volin, E. Dereniak, T. Gleeson, M. Hopkins, D. Wilson, and P. Maker, "Demonstration of a computed-tomography imaging spectrometer using a computer-generated hologram disperser," *Appl. Opt.* **36**, 3694–3698 (1997).
9. M. E. Gehm and D. J. Brady, "High-throughput hyperspectral microscopy," *Proc. SPIE* **6090**, 13–21 (2006).
10. M. Gehm, S. McCain, N. Pitsianis, D. Brady, P. Potluri, and M. Sullivan, "Static two-dimensional aperture coding for multimodal, multiplex spectroscopy," *Appl. Opt.* **45**, 2965–2974 (2006).
11. S. McCain, M. Gehm, Y. Wang, N. Pitsianis, and D. Brady, "Coded Aperture Raman Spectroscopy for Quantitative Measurements of Ethanol in a Tissue Phantom," *Appl. Spectrosc.* **60**, 663–671 (2006).

12. E. Cull, M. Gehm, D. Brady, C. Hsieh, O. Momtahan, and A. Adibi, "Dispersion multiplexing with broadband filtering for miniature spectrometers," *Appl. Opt.* **46**, 365–374 (2007).
13. C. Fernandez, B. Guenther, M. Gehm, D. Brady, and M. Sullivan, "Longwave infrared (LWIR) coded aperture dispersive spectrometer," *Opt. Express* **15**, 5742–5753 (2007).
14. M. E. Gehm, D. J. Brady, N. Pitsianis, and X. Sun, "Compressive sampling strategies for integrated microspectrometers," in *Proc. SPIE*, R. A. Athale and J. C. Zolper, eds., vol. 6232 (SPIE, 2006).
15. D. J. Brady and M. E. Gehm, "Compressive imaging spectrometers using coded apertures," in *Proc. SPIE*, Z. ur Rahman, S. E. Reichenbach, and M. A. Neifeld, eds., vol. 6246 (SPIE, 2006).
16. R. M. Willett, M. E. Gehm, and D. J. Brady, "Multiscale reconstruction for computational spectral imaging," in *Proc. SPIE*, I. P. Charles A. Bouman, Eric L. Miller, ed., vol. 6498 (SPIE-IS and Electronic Imaging, 2007).
17. E. Candès, J. Romberg, and T. Tao, "Robust Uncertainty Principles: Exact Signal Reconstruction from Highly Incomplete Frequency Information," *IEEE Trans. Inf. Theory* **52**, 489 – 509 (2006).
18. E. Candès and T. Tao, "Near Optimal Signal Recovery From Random Projections: Universal Encoding Strategies," (2006). To be published in *IEEE Trans. Inf. Theory*. <http://www.acm.caltech.edu/emmanuel/papers/OptimalRecovery.pdf>.
19. D. Donoho, "Compressed Sensing," *IEEE Trans. Inf. Theory* **52**, 1289–1306 (2006).
20. M. Harwit and N. Sloane, *Hadamard Transform Optics* (Academic Press, New York, 1979).
21. E. Candès, J. Romberg, and T. Tao, "Stable Signal Recovery from Incomplete and Inaccurate Measurements," *Commun. Pure Appl. Math.* **59**, 1207–1223 (2006).
22. J. Haupt and R. Nowak, "Signal Reconstruction from Noisy Random Projections," (2006). To be published in *IEEE Trans. Inf. Theory*. <http://www.ece.wisc.edu/nowak/infth.pdf>.
23. W. Richardson, "Bayesian-based iterative method of image restoration," *J. Opt. Soc. of Am.* **62**, 55–59 (1972).
24. L. B. Lucy, "An iterative technique for the rectification of observed distributions," *Astron. J.* **79**, 745–754 (1974).
25. R. Nowak and E. Kolaczyk, "A Multiscale Statistical Framework for Poisson Inverse Problems," *IEEE Trans. Inf. Theory* **46**, 1811–1825 (2000).
26. R. Willett, "Multiscale intensity estimation for marked Poisson processes," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)* (2007).
27. E. Kolaczyk and R. Nowak, "Multiscale Likelihood Analysis and Complexity Penalized Estimation," *Annals of Stat.* **32**, 500–527 (2004).
28. R. Willett and R. Nowak, "Fast Multiresolution Photon-Limited Image Reconstruction," in *Proc. IEEE Int. Sym. Biomedical Imaging — ISBI '04* (15-18 April, Arlington, VA, USA, 2004).
29. M. Lang, H. Guo, J. E. Odegard, C. S. Burrus, and R. O. Wells, "Noise reduction using an undecimated discrete wavelet transform," *IEEE Signal Processing Lett.* **3**, 10–12 (1996).

1. Introduction

Spectral imaging is an emerging tool for a variety of scientific and engineering applications because of the additional information it provides about the nature of the materials being imaged. Traditional imagers produce two-dimensional spatial arrays of scalar values representing the intensity of a scene. A spectral imager, in contrast, produces a two-dimensional spatial array of vectors which contain the spectral information for the respective spatial locations. The resulting data is known as the *spectral data cube* because of its three-dimensional nature. The addition of spectral information can provide valuable information in a variety of contexts ranging from environmental monitoring [1, 2], to astrophysics [3], to biochemistry [4, 5] to security applications [6].

Adoption of spectral imaging has been slow because of a fundamental tradeoff between spatial resolution, spectral resolution, light collection, and measurement acquisition time. Standard spectral imaging designs can simultaneously optimize only two of the four quantities—resulting in relatively poor overall performance. The origin of these tradeoffs can be readily understood. Traditional, non-multiplexed, spectrometers already exhibit a tradeoff between light collection and spectral resolution. Any system that attempts to use one of these systems as the spectrograph in a spectral imager inherits this limitation. Further, the fact that the spectral image data cube is three-dimensional, while available detector arrays are two-dimensional results in either a need for scanning (which increases the overall acquisition time) or in the tiling of the detector array with multiple two-dimensional slices of the cube (which, if the field of view is held fixed, limits the spatial resolution). In the past decade, however, there have been a number of

ingenious designs that allow independent control of three of these quantities simultaneously [7, 8, 9], largely through the introduction of various multiplex measurement techniques.

This paper attempts to eliminate the remaining tradeoffs by presenting a novel spectral imaging system and associated reconstruction framework that fully decouples the four operational quantities. The imager is a completely static, single-shot design, resulting in a mechanically robust and inexpensive system. In these respects the system architecture is a descendent of the coded-aperture spectroscopy architecture previously developed by several of the current authors [10, 11, 12, 13, 14]. In our first extension to spectral imaging [9], however, that implementation required a sequence of exposures in order to measure the contents of the spectral data cube. The system architecture described in this manuscript (of which early versions were discussed in [15, 16]) has been developed precisely to avoid that shortcoming. In this system, the imager does not directly measure each voxel in the data cube. Instead, it collects a small number (relative to the size of the data cube) of coded measurements, and then a novel reconstruction method is used to estimate the spectral image from the noisy projections.

This approach draws heavily on ideas in the emerging field of compressed sensing [17, 18, 19]. In compressed sensing, certain design strategies are incorporated into measurement systems in a way that can dramatically improve the system's ability to produce high-quality reconstructions from a limited number of measurements. The basic idea of this theory is that when the signal of interest is very sparse (ie. zero-valued at most locations) or highly compressible in some basis, relatively few *incoherent* observations are necessary to reconstruct the most significant non-zero signal components. In the remainder of this manuscript we demonstrate the practical application of these ideas to spectral imaging. We describe a particular system design, present a multiscale reconstruction algorithm, and demonstrate that accurate spectroscopic images can be estimated from an under-determined set of noisy projections.

2. System design

The system is comprised of two sequential dispersive arms of the 4-f type commonly used (singly) as a traditional dispersive spectrometer. The two arms are arranged in opposition so that the second arm exactly cancels the dispersion introduced by the first arm. A coding aperture occupies the plane separating the two arms. A schematic of the system is shown in Fig. 1.




Fig. 1. Schematic of the spectral imager.

The operational characteristics of the system can be easily understood on a conceptual level. A standard imaging relay (not shown) is used to form an image of a remote scene in the plane of the input aperture. The input aperture is then imaged through the first arm onto the plane containing the coding aperture. However, because the arm contains a dispersive element, multiple images are formed at wavelength-dependent locations. At this point, the spatial structure

in the plane of the coding aperture contains a mixture of spatial and spectral information about the source. Passing through the coding aperture modulates this information with the applied pattern. The second arm then undoes the spatial-spectral mixing introduced by the first arm as it forms an image of the source on the detector. In the process of undoing the effects of arm 1, the spatial modulation introduced by the coding aperture is transformed into a spatial and spectral modulation. The details of the process are described in the following section.

2.1. System model

We denote the spectral density entering the instrument as $S_0(x, y; \lambda)$. The spectral density just prior to the code aperture is then

$$\begin{aligned} S_1(x, y; \lambda) &= \iint dx' dy' \delta(x' - [x + \alpha(\lambda - \lambda_c)]) \delta(y' - y) S_0(x', y'; \lambda) \\ &= S_0(x + \alpha(\lambda - \lambda_c), y; \lambda). \end{aligned} \quad (1)$$

Here the Dirac delta functions describe propagation through unity-magnification imaging optics and a dispersive element with linear dispersion α and center wavelength λ_c . (Note that this model assumes linear dispersion—which is approximately true only over limited wavelength ranges. The system can still be operated in nonlinear regions, as the processing algorithm is calibration based. The linear model is only used to provide guidance about the aperture code.)

Immediately after the coding aperture the spectral density is given by

$$S_2(x, y; \lambda) = T(x, y) S_1(x, y; \lambda) = T(x, y) S_0(x + \alpha(\lambda - \lambda_c), y; \lambda), \quad (2)$$

with $T(x, y)$ the spatial transmission pattern imposed by the coding aperture.

Propagation through the second set of imaging optics and the second dispersive element results in a spectral density in the detector plane of

$$\begin{aligned} S_3(x, y; \lambda) &= \iint dx' dy' \delta(x' - [x - \alpha(\lambda - \lambda_c)]) \delta(y' - y) S_2(x', y'; \lambda) \\ &= T(x - \alpha(\lambda - \lambda_c), y) S_0(x, y; \lambda) \\ &= H(x, y; \lambda) S_0(x, y; \lambda). \end{aligned} \quad (3)$$

Again, the Dirac delta functions describe propagation through the imaging optics and disperser (note the internal sign change representing the reversed orientation of the disperser). Here we have defined the spectral density filter function $H(x, y; \lambda) = T(x - \alpha(\lambda - \lambda_c), y)$. This formulation makes explicit the fact that the two-dimensional coding pattern introduces a three-dimensional filter function that acts on the source spectral density.

As the detector is wavelength-insensitive, the ultimate quantity measured is not the spectral density in the detector plane, but the intensity

$$I(x, y) = \int d\lambda H(x, y; \lambda) S_0(x, y; \lambda). \quad (4)$$

Further, the detector array is spatially pixelated. If we take the pixel size as Δ , then the detector measurements become

$$I_{nm} = \iiint dx dy d\lambda \text{rect}\left(\frac{x}{\Delta} - m, \frac{y}{\Delta} - n\right) H(x, y; \lambda) S_0(x, y; \lambda). \quad (5)$$

It then becomes natural to consider a coding aperture where the transmission function T is pixelated with features that are the same size as the detector pixels (Note that for implementation reasons, it is actually preferable to consider codes where the features sizes are integer

multiples of the detector pixel size. Extending this formal treatment to that case is straightforward.)

$$T(x, y) = \sum_{m', n'} T_{n'm'} \text{rect} \left(\frac{x}{\Delta} - m', \frac{y}{\Delta} - n' \right). \quad (6)$$

With this, Eq. 5 becomes

$$I_{nm} = \sum_{m', n'} \iiint dx dy d\lambda \text{rect} \left(\frac{x}{\Delta} - m, \frac{y}{\Delta} - n \right) \text{rect} \left(\frac{x - \alpha(\lambda - \lambda_c)}{\Delta} - m', \frac{y}{\Delta} - n' \right) \times T_{n'm'} S_0(x, y; \lambda). \quad (7)$$

To gain an understanding of how the coding pattern influences the measured intensity distribution, it is instructive to consider two highly constrained cases. First we consider a monochromatic source with $S_0(x, y, \lambda) = I_0(x, y) \delta(\lambda - \lambda_c)$, where $I_0(x, y)$ is the intensity distribution of the monochromatic scene. In this case, Eq. 7 simplifies significantly

$$\begin{aligned} I_{nm}|_{\lambda=\lambda_c} &= \sum_{m', n'} \iiint dx dy d\lambda \text{rect} \left(\frac{x}{\Delta} - m, \frac{y}{\Delta} - n \right) \text{rect} \left(\frac{x}{\Delta} - m', \frac{y}{\Delta} - n' \right) \\ &\quad \times T_{n'm'} I_0(x, y) \delta(\lambda - \lambda_c) \\ &= \sum_{m', n'} \delta_{mm'} \delta_{nn'} T_{n'm'} I_{0, nm} \\ &= T_{nm} I_{nm}, \end{aligned} \quad (8)$$

where $I_{0, nm}$ is a spatially-pixelated version of the monochromatic source intensity $I_0(x, y)$. Next, we consider the response to a monochromatic source at $\lambda = \lambda_c + \Delta\lambda$ with $\Delta\lambda = \Delta/\alpha$, with α the linear dispersion of the dispersive elements. In this case, we find

$$\begin{aligned} I_{nm}|_{\lambda=\lambda_c+\Delta\lambda} &= \sum_{m', n'} \iiint dx dy d\lambda \text{rect} \left(\frac{x}{\Delta} - m, \frac{y}{\Delta} - n \right) \text{rect} \left(\frac{x}{\Delta} - (m' + 1), \frac{y}{\Delta} - n' \right) \\ &\quad \times T_{n'm'} I_0(x, y) \delta(\lambda - (\lambda_c + \Delta\lambda)) \\ &= \sum_{m', n'} \delta_{mm'} \delta_{nn'} T_{n'(m'+1)} I_{0, nm} \\ &= T_{n(m-1)} I_{0, nm}. \end{aligned} \quad (9)$$

Thus we see that the contribution from a particular wavelength is the pixelated source weighted by a version of the aperture code that shifts in the x direction as the wavelength changes. At the center wavelength, the weighting pattern is aligned with the detector pixels. For wavelength shifts that are integer multiples of $\Delta\lambda$, the pattern is again registered with the detector pixels. However, for intermediate wavelengths, the elements in the coding pattern straddle multiple detector pixels.

Finally, we define the filter function

$$\begin{aligned} w_{nmp} &= \sum_{m', n'} \iiint dx dy d\lambda \text{rect} \left(\frac{x}{\Delta} - m, \frac{y}{\Delta} - n, \frac{\lambda - \lambda_c}{\Delta\lambda} - p \right) \\ &\quad \times \text{rect} \left(\frac{x - \alpha(\lambda - \lambda_c)}{\Delta} - m', \frac{y}{\Delta} - n' \right) T_{n'm'}, \end{aligned} \quad (10)$$

and the fully pixelated source spectral density as

$$s_{nmp} = \frac{1}{\Delta^2 \Delta\lambda} \iiint dx dy d\lambda \text{rect} \left(\frac{x}{\Delta} - m, \frac{y}{\Delta} - n, \frac{\lambda - \lambda_c}{\Delta\lambda} - p \right) S_0(x, y; \lambda). \quad (11)$$

If the source spectrum is slowly varying on the scale $\Delta\lambda$, then we can approximate Eq. 5 as

$$I_{nm} = \sum_p w_{nmp} s_{nmp}. \quad (12)$$

2.2. Code design

The analysis in the previous section determined the nature of the spatio-spectral filter that is applied and how it is related to the code pattern placed in the system. It is most convenient to view the filter function w_{nmp} as a spatially varying collection of spectral filters. In this framework, the analysis above reveals the following facts:

1. The spectral filter for a given spatial location may contain structure on a scale as narrow as $\Delta\lambda = \Delta/\alpha$. Typically, this is significantly narrower than the features in a traditional thin-film interference filter.
2. The spectral filters in different rows of the detector (different n) are unconstrained. There are no correlations imposed by the physical structure.
3. The spectral filters within a row (same n) are cyclic permutations of each other. This arises from the spatial shift in the m -direction that arises from spectral shifts. (e.g., In a hypothetical instrument with 4 spectral channels, the spectral filter applied to a particular pixel had values $\{a, b, c, d\}$, then the filters on its left/right neighbors are $\{b, c, d, a\}$ and $\{d, a, b, c\}$, respectively).
4. The measurement at a given detector pixel is the inner product of the spectral filter for that pixel and the pixelated source spectrum for that spatial location (see Eq. 12).

To create a system with M spectral channels requires a 1-D code of at least length M . For the remainder of this paper, we assume a code of length M for simplicity. The physical nature of the system produces spectral filters that are all M possible cyclic shifts of the fundamental code. Thus we are led to consider well-conditioned codes that consist of cyclic permutations of a single master codeword. The canonical example of these types of codes are those based on the *cyclic S-matrices* [20]. For our initial system (and earliest simulations), we drew upon the order-15 cyclic S-matrix with the fundamental codeword “100100011110101”. This code provides several advantages, including M *unique* cyclic shifts and an overall transmission efficiency of $8/15$.

If the coding plane were directly tiled with this pattern, the various filter functions would be implemented on the detector plane in the manner depicted in Fig. 2. In this image, the number k denotes the k^{th} spectral filter function, which corresponds to the fundamental codeword circularly shifted by $k - 1$ bits; thus 1 refers to the fundamental codeword “100100011110101”, 2 refers to the shifted codeword “001000111101011”, etc. The difficulty with this arrangement is that there are no compact regions that contain all of the filter functions.

15	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	1
15	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	1
15	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	1

Fig. 2. Distribution of filter functions that arises from simple tiling of the fundamental codeword. No compact region contains all 15 filters.

If we instead tile the aperture plane with a unit cell of the form

```

100100011110101
001111010110010
101011001000111

```

(three copies of the fundamental codeword, but with circular shifts of five elements between rows) then the filter functions are arrayed on the detector plane as displayed in Fig. 3. In this scheme, any 3×5 region on the detector contains all 15 spectral filters.

10	11	12	13	14	15	1
15	1	2	3	4	5	6
5	6	7	8	9	10	11
10	11	12	13	14	15	1
15	1	2	3	4	5	6

Fig. 3. Distribution of filter functions that arises from the more complicated unit cell. Any 3×5 region contains all 15 filters.

3. Reconstruction methods

In this section we describe the multiscale reconstruction process employed to get the spatial and spectral information from the mask-modulated intensity information represented in Eq. 12. First we formulate a parametric model for the noisy observations we measure at the detector array, and then describe the estimation process. We compute an optimal solution to this underdetermined problem using an expectation maximization algorithm combined with a wavelet denoising technique. The following subsections give a systematic description of our approach.

3.1. Multiscale reconstruction method

Let the spectral image of interest be denoted $\mathbf{s} = \{s_{nmp}\}$, and the intensity of the observations as $\mathbf{I} = \{I_{nm}\}$, so that equation (12) can be written in matrix-vector notation as

$$\mathbf{I} = \mathbf{W}\mathbf{s},$$

where the matrix \mathbf{W} performs the discretized filtering described in Sec. 2.1. In addition, we model our observed data as

$$\mathbf{d} \sim \text{Poisson}(\mathbf{I}) = \text{Poisson}(\mathbf{W}\mathbf{s}),$$

so that the likelihood of observing \mathbf{d} given spectral image \mathbf{s} is

$$p(\mathbf{d}|\mathbf{W}\mathbf{s}) = \prod_{n=1}^N \prod_{m=1}^M \frac{e^{-\sum_p w_{nmp}s_{nmp}} (\sum_p w_{nmp}s_{nmp})^{d_{nm}}}{d_{nm}!}.$$

Note that \mathbf{W} has many more columns than rows (by a factor of M), making this a very underdetermined problem.

To solve this challenging inverse problem, we seek a solution ($\hat{\mathbf{s}}$) which is both a good match to the data (\mathbf{d}) and *sparse*. In particular, we solve the following optimization problem:

$$\hat{\mathbf{s}} = \arg \min_{\tilde{\mathbf{s}} \in \mathcal{S}} \{-\log p(\mathbf{d}|\mathbf{W}\tilde{\mathbf{s}}) + \text{pen}(\tilde{\mathbf{s}})\}, \quad (13)$$

where \mathcal{S} is a collection of estimators to be described below, and $\text{pen}(\tilde{\mathbf{s}})$ is a penalty is proportional to the sparsity of $\tilde{\mathbf{s}}$ in a multiscale spatio-spectral representation. This will be described in detail below.

The use of sparsity to solve challenging and ill-posed inverse problems has received widespread attention recently [17, 18, 19, 21, 22]. The objective function proposed in (13) in particular is similar to those proposed in [21, 22], in that we seek a solution accurately represented by a small number of cells in an RDP (i.e. sparse). We compute the solution to this problem using an Expectation-Maximization algorithm, which in this case is a regularized version of the Richardson-Lucy algorithm [23, 24, 25]. The method consists of two alternating steps:

Step 1:

$$\mathbf{y}^{(t)} = \hat{\mathbf{s}}^{(t)} \cdot \times \mathbf{W}^T(\mathbf{d} / \mathbf{W}\hat{\mathbf{s}}^{(t)}),$$

where $\cdot \times$ and $\cdot /$ denote element-wise multiplication and division, respectively.

Step 2: Compute $\hat{\mathbf{s}}^{(t+1)}$ by denoising $\mathbf{y}^{(t)}$:

$$\hat{\mathbf{s}}^{(t+1)} = \arg \min_{\tilde{\mathbf{s}}} \left\{ -\log p(\mathbf{y}^{(t)} | \tilde{\mathbf{s}}) + \text{pen}(\tilde{\mathbf{s}}) \right\}.$$

The penalized likelihood denoising method employed in Step 2 takes advantage of correlations in the data between both wavelengths and spatial locations. The proposed method entails performing hereditary Haar intensity estimation via tree pruning in the spatial dimensions, with each leaf of the resulting unbalanced quad-tree decomposition corresponding to a region of spatially homogeneous spectra.

In particular, we determine the ideal partition of the spatial domain of observations and use maximum likelihood estimation to fit a single spectrum to each square in the optimal spatial partition. The space of possible partitions is a nested hierarchy defined through a recursive dyadic partition (RDP) of the datacube domain. The optimal partition is selected by merging neighboring squares of (i.e. pruning) a quad-tree representation of the observed data to form a data-adaptive RDP \mathcal{P} . Each of the terminal squares in the pruned spatial RDP could correspond to a region of intensity which is spatially homogeneous or smoothly varying (regardless of the regularity or irregularity between the spectral bands). This gives our estimators the capability of spatially varying the resolution to automatically increase the smoothing in very regular regions of the intensity and to preserve detailed structure in less regular regions.

Given a partition \mathcal{P} , $\tilde{\mathbf{s}}(\mathcal{P})$ can be calculated by finding the “best” spectrum fit to the observations over each cell in \mathcal{P} . This can be accomplished by simply computing the mean observed spectrum in each cell of \mathcal{P} . The final spatio-spectral estimate is then calculated by finding the partition which minimizes the total penalized likelihood function:

$$\begin{aligned} \hat{\mathcal{P}} &\equiv \arg \min_{\mathcal{P}} \left\{ -\log p(\mathbf{y}^{(t)} | \tilde{\mathbf{s}}(\mathcal{P})) + \text{pen}(\mathcal{P}) \right\} \\ \hat{\mathbf{s}} &\equiv \tilde{\mathbf{s}}(\hat{\mathcal{P}}), \end{aligned} \quad (14)$$

where $\text{pen}(\mathcal{P})$ is a penalty proportional to the number of cells in the RDP, encouraging a sparse solution (in terms of the size of the RDP). This method is described in detail in [26].

This approach is similar to the image estimation method described in [27, 28], with the key distinction that the proposed method *forces the spatial RDP to be the same at every spectral band*. A sample such partition is displayed in Fig. 4. This constraint makes it impossible for the method to perform spatial smoothing at some spectral bands but not others. In other words, when a tree branch is pruned in the proposed framework, it means partition cells are merged

in every spectral band simultaneously at the corresponding spatial location. This approach is effective because an outlier observation in one spatio-spectral voxel may not be recognized as such when spectral bands are considered independently, but may be correctly pruned when the corresponding spectrum is very similar to spatially nearby spectra.




Fig. 4. Sample partition of a spatio-spectral data cube. The spatial partition is the same at each spectral band, making it impossible for the estimation method to perform spatial smoothing at some spectral bands but not others.

The accuracy of these estimates can be augmented by a process called cycle-spinning, or averaging over shifts, resulting in translation-invariant (TI) estimates [29]. Cycle-spinning was derived in the context of undecimated wavelet coefficient thresholding in the presence of Gaussian noise, but is difficult to implement efficiently in the case of tree-pruning methods. The above multiscale tree-pruning methods can be modified to produce the same effect by averaging over shifts, but the increase in quality comes at a high computational cost. Novel computational methods [28], however, can be used to yield TI-Haar tree pruning estimates in $O(N^2 M \log N \log M)$ time for an $N \times N \times M$ data cube.

4. Experimental results

To test the dual-disperser architecture, we constructed a proof-of-concept spectral imaging system as described in Sec. 2. A photograph of the experimental prototype is shown in Fig. 5. As designed and built, the prototype has an operational range of 520–590 nm. It is an interesting consequence of the system approach that shifting the wavelength by one spectral band should produce the same physical shift on the coding array as a spatial shift by one spatial resolution element. As a result, the system requires a low amount of dispersion. For this prototype we achieved this with a prism-based approach (diffractive approaches are possible, but more challenging as low dispersion tends to introduce overlap from higher diffraction orders).

Before the system may be used to for spectral imaging, it requires calibration. This calibration is performed by sequentially measuring the response of the system to spatially-uniform monochromatic light at 2 nm steps in the range 520–590 nm. To minimize the noise in the calibration, the system response at a given wavelength is taken to be the average of 10 exposures. Further, the monochromator used to generate the light has a non-uniform output as a function of wavelength. To control for this, a small portion of the light output is fed into a photodiode with a known wavelength response. At each wavelength step, the output of this photodiode is used to adjust the exposure time to keep the total energy propagating through the spectral im-




Fig. 5. The experimental prototype.

ager constant. This process is automated via MATLAB. When the resulting system responses are stacked according to wavelength, the result is the system filter function w_{nmp} as defined in Eq. 10

As a first test of the system we consider simple targets illuminated with monochromatic light. We begin with a pingpong ball and illuminations at 532 nm and 543 nm. A representative detector image (taken with 532 nm illumination) is shown in Fig. 6(a). Note the modulation introduced by the coding aperture. The pattern is particularly clear because of the monochromatic source. An estimated intensity image of the scene, generated by summing the spectral channels in the reconstructed datacube is shown in Fig. 6(b). The reconstruction algorithm has largely succeeded in eliminating the modulation introduced by the screen. Figures 6(b) and (c) are the spectral estimates for a particular spatial location in the cases with 532 nm and 543 nm illumination, respectively. The reconstruction accurately locates the spectral peaks.

We next consider the same simple target but now with slightly broader illumination—broadband light filtered by 10 nm FWHM bandpass filters centered at 560 nm and 580 nm, respectively. A representative detector image take with the 560 nm filter is shown in Fig. 7(a). Note that the sharpness of the modulation pattern introduced by the coding aperture has decreased as a result of the increased spectral width. An estimated intensity image of the scene is shown in Fig. 7(b). Figures 7(c) and (d) show the spectral reconstruction at a particular spatial location for the 560 nm and 580 nm bandpass filters, respectively. Again, the reconstruction accurately locates the spectral peaks. However, we note that, for the 580 nm bandpass, a spurious peak now appears near 520 nm. For this bandpass, there is significant source illumination at wavelengths longer than 590 nm. The system is designed for a spectral range of 520–590 nm. Any wavelengths outside this band are aliased back into this range (the illumination produces a coding pattern that is indistinguishable from that generated by a wavelength inside the operational band). The peak at 520 nm, then, is an aliased version of source energy above 590 nm. Any final system would avoid this difficulty by incorporating a bandpass filter matched to the spectral range of the instrument.

Finally, we image a slightly less constrained target under more normal illumination. For this portion of the experiment, we use a collection of citrus fruit under broadband (white) light illumination. The resulting reflection spectrum contains significant energy in the spectral band of interest. Figure 8(a) shows the detector image captured during this experiment. The top fruit is green, the bottom right is yellow-green, and the bottom left is yellow-orange. Note the broad spectral ranges have made the spectral patterns very indistinct, and that the patterns are subtly different in the three regions as a result of the different spectral components. A reconstructed intensity image is shown in Fig. 8(b). The spectral reconstructions for the three different regions

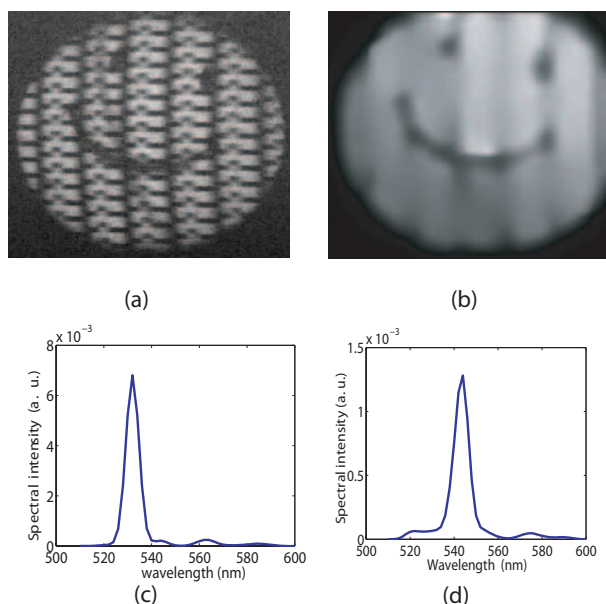


Fig. 6. Experimental results from simple targets with monochromatic illumination. (a) Detector image recorded for 532 nm illumination. (b) Intensity image generated by summing the spectral information in the reconstruction for 532 nm illumination. (c) Spectral reconstruction at a particular spatial location for 532 nm illumination. (d) Spectral reconstruction at a particular spatial location for 543 nm illumination.

are shown in Fig. 8(c) along with measurements made by a conventional spectrometer. There is good qualitative agreement between the reconstructions and the conventional measurements, especially in the central spectral regions. We believe that the measurements from the conventional instrument are not fully accurate because of the practical difficulties in coupling the reflected spectra to the probe. However, at least a portion of the deviation arises from inaccurate reconstructions, as we again see some spurious peaks near 520 nm that are the result of spectral aliasing.

Figure 9 shows slices of the reconstructed datacube for 8 specific channels across the spectral range. An animation showing all of the channels in the reconstruction is in Fig. 10.

5. Conclusions

In this manuscript we have described a new, single-shot spectral imager based on compressive sensing ideas. The ability to create nearly arbitrary spectral projections on the source datacube arises from the novel dual-disperser design. In addition, we have developed a unique multiscale reconstruction method for use with compressive spectral imagers of this type. This reconstruction technique combines a maximum likelihood estimator with a penalty-based multiscale denoising technique that utilizes spatio-spectral correlations in the scene to improve the reconstruction quality.

A proof-of-concept prototype has been constructed and tested on both highly-constrained and real-world sources. The reconstructions accurately capture the spectral features of the source (barring a spectral aliasing that can be eliminated through the incorporation of a bandpass filter on future systems). The spatial structure of the reconstructions is also accurate. The modulation introduced by the coding aperture is successfully removed by the reconstruction (especially for




Fig. 7. Experimental results from simple targets with narrow-band illumination. (a) Detector image recorded for illumination with a 10 nm FWHM bandpass centered at 560 nm. (b) Intensity image generated by summing the spectral information in the reconstruction for the 560 nm bandpass. (c) Spectral reconstruction at a particular spatial location for the 560 nm bandpass. (d) Spectral reconstruction at a particular spatial location for the 580 nm bandpass. The origin of the small peak near 520 nm is explained in the text.

the broadband, real-world scene).

The result is a single-shot spectral imager that, for the first time, mitigates the trade-offs between spatial resolution, spectral resolution, light collection, and measurement acquisition time. While the performance of the system is quite acceptable for a first proof-of-concept, we feel that the results are limited by the stock optics used to create the present prototype. We are currently constructing a new version, with dramatically enhanced spatial and spectral resolution.

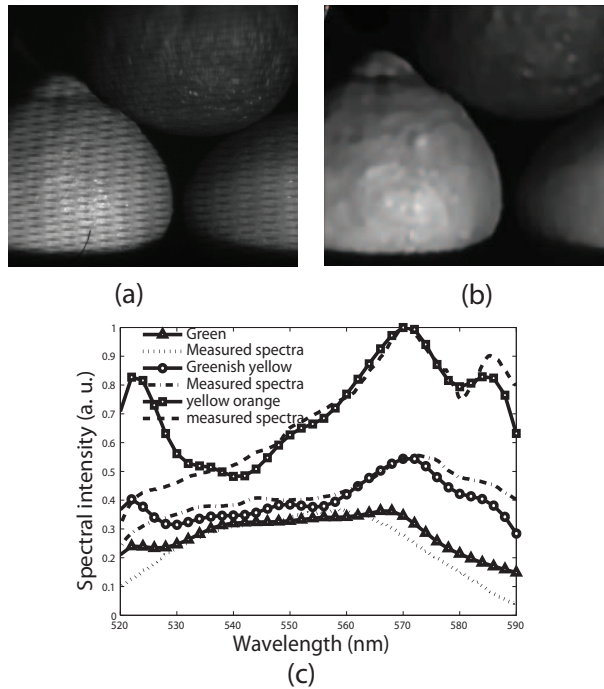


Fig. 8. Experimental results from real-world objects under broadband (white) illumination. (a) Detector image recorded by the system. (b) Reconstructed intensity image of the scene. (c) Spectral reconstructions for spatial locations in the three regions.




Fig. 9. Slices through the reconstructed datacube at 8 particular spectral channels.




Fig. 10. Animation showing all 36 reconstructed spectral channels. (Multimedia file, 1.1 MB.)